Volume 31, Number 1, 1987

Digital halftoning on the IBM 4250 Printer *by G. Goertzel and G. R. Thompson, p. 2.* A method of reproducing high-quality continuous-tone images via the IBM 4250 Printer is presented. The approach is modeled on the halftone process used in conventional lithography but is adapted to discrete bilevel printing and digital processing. We use a combination of standard and novel techniques. These include generation and calibration of a set of halftone dot patterns, randomized error propagation in pattern selection, and resolution enhancement in areas of high intensity gradients. The resultant images have good gray-scale rendition and sharp edges without the problems of contouring and worminess often associated with digital image reproduction. We have named this approach PREPRESS: Picture Rendition using Error Propagation and Resolution Enhancement in Simulated Screening.

Binary-image-manipulation algorithms in the Image View Facility *by K. L. Anderson, F. C. Mintzer, G. Goertzel, J. L. Mitchell, K. S. Pennington, and W. B. Pennebaker, p. 16.* Most current implementations of electronic mail deal primarily with coded information. A scanned-document-handling system that could scan a document, distribute it, display it on terminals, and print it on host-attached printers would offer a similar convenience for documents in hard-copy rather than coded form. For such a system to be practical, fast software is needed for a number of image-manipulation functions. The required functions are compression, to reduce the size of the data files; decompression, to reconstruct the scanned document; scaling, to match the resolution of the scan to the resolution of the display or printer; and rotation, to reorient documents scanned sideways or upside down. This paper describes a collection of algorithms underlying fast software for manipulating binary images that is used in the Image View Facility, a System/370-based software package that permits the display and printing of binary images at various resolutions.

PANDA: Processing Algorithm for Noncoded Document Acquisition *by Y.-H. Chen, F. C. Mintzer, and K. S. Pennington, p. 32.* With a scanned-document-handling system, documents can be scanned, stored, transmitted to remote locations, viewed on displays and terminals, edited, and printed. These systems hold much promise for office automation, since they facilitate the communication and storage of information that is not easily recoded into the traditional formats for text and graphics. However, most of the current systems that perform these functions are intended for documents that at every point are either black or white, but not gray. These systems effectively exclude documents that contain regions with varying shades of gray (known as continuous-tone regions). PANDA, the Processing Algorithm for Noncoded Document Acquisition, is a technique that processes *mixed documents,* those that contain continuous-tone regions in addition to text, graphics, and line art. PANDA produces a high-quality binary representation of all regions of a mixed document. Furthermore, all regions of the binary representation, including the continuous-tone image regions, are significantly compressed by the run-length-based compression algorithms that underlie scanned-document-handling systems. This is a key feature of PANDA. Indeed, this compression makes practical the inclusion of mixed documents into many existing systems.

YODA: An advanced display for personal computers *by S. Gupta, D. F. Bantz, P. N. Sholtz, C. J. Evangelisti, and W. R. DeOrazio, p. 44.* YODA (the YOrktown Display Adapter) is an experimental display designed to improve the quality and speed of users' interactions with personal computers. This paper describes the YODA hardware architecture and software design. Special attention is given to techniques used for antialiasing. The various trade-offs and decisions that were made are discussed.

Document convergence in an interactive formatting system *by D. D. Chamberlin, p. 58.* One of the most complex aspects of document formatting is the processing of references to remote objects such as headings and figures. In the case of a forward reference to an object that occurs later in the document, two formatting passes are usually needed before the document converges to a stable state. Some documents require more than two passes to converge, and cases are known of documents that never converge but oscillate between two unstable states. This paper describes the techniques used for resolving references and detecting document convergence by the Interactive Composition and Editing Facility, Version 2 (ICEF2). ICEF2 is an interactive formatting system that allows users to move about in a document, editing and reformatting pages. The concepts of *formatting pass* and *document convergence* are discussed in the context of interactive formatting. A description is given of the ICEF2 data store, a small relational database manager with special features for detecting document convergence. A sample ICEF2 style definition is discussed to illustrate how ICEF2 deals with document elements whose appearance depends on their location on the page.

LEXX—A programmable structured editor *by M. F. Cowlishaw, p. 73.* Many sophisticated and specialized editing programs have been developed in recent years. These editors help people manipulate data, but the diversity complicates rather than simplifies computer use. LEXX is an editing program that can work with the syntax and structure of the data it is presenting, yet is not restricted to just one kind of data. It is used for editing programs, documents, and other material and hence provides a consistent environment for the user regardless of the editing task. The new *live parsing* technique used by LEXX means that it can be programmed to handle a very wide variety of structured data. The structure information is, in turn, used to improve the presentation of data (through color, fonts, and formatting), which makes it easier for people to deal with the text being edited.

A method for efficient storage and rapid application of context-sensitive phonological rules for automatic speech recognition *by R. L. Mercer and P. S. Cohen, p. 81.* In an automatic speech-recognition system, the application of phonological rules to phonemic strings in order to create phonetic graphs is a computationally time- and storage-consuming process. A great many such graphs must be constructed during the decoding phase; thus it is important to be able to rapidly construct phonetic graphs for strings of words from graphs of individual words. However, because many phonological rules operate across word boundaries or require interword context, it is not possible to determine a unique, context-independent phonetic graph for a word. We describe a method for determining the phonetic graph for a word in isolation, together with auxiliary information to allow phonetic graphs for different words to be rapidly interconnected to form a phonetic graph for a string of words; the method also reduces storage requirements significantly.

Feature analysis for symbol recognition by elastic matching *by J. M. Kurtzberg, p. 91.* A technique has been developed for the recognition of unconstrained handwritten discrete symbols based on elastic matching against a set of prototypes generated by individual writers. The incorporation of feature analysis with elastic matching to eliminate unlikely prototypes is presented in this paper and is shown to greatly reduce the required processing time without any deterioration in recognition performance.

Drop formation by DOD ink-jet nozzles: A comparison of experiment and numerical simulation *by T. W. Shield, D. B. Bogy, and F. E. Talke, p. 96.* This paper presents a comparison of a numerical simulation of drop formation and ejection from a drop-on-demand (DOD) ink-jet nozzle with experimental observations from a particular nozzle-transducer design. In the numerical simulation, first the pressure waves in the transducer chamber are calculated using inviscid compressible flow theory to obtain the pressure history at the inner face of the nozzle plate. Then a viscous momentum integral computation is applied to the nozzle to obtain the velocity history at the outer face of the nozzle plate. Finally, the free surface shape is calculated using finite-difference methods on the one-dimensional equations for an inviscid incompressible free jet with surface tension that uses the nozzle exit velocity history as the driving boundary condition. The computations are compared with drop formation photographs obtained from a particular nozzle-transducer design. Encouraging agreement is obtained, but the numerical model will require added sophistication before detailed agreement can be expected.

DNA vectorgrams: Representation of cancer genes as movements on a 2D cellular lattice *by C. A. Pickover, p. 111.* A brief introduction introduction to a computer graphics characterization of cancer DNA sequences, as well as other biologically interesting sequences, is presented. The procedure described takes DNA sequences containing n bases and computes n two-dimensional real vectors. When displayed on a planar unit-cellular lattice, these characteristic patterns appear as a "DNA vectorgram," $C(n)$. Several demonstration plots are provided which indicate that $C(n)$ is sensitive to certain statistical properties of the sequence of bases and allows the human observer to visually detect some important sequence structural properties and patterns not easily captured by traditional methods. The system presented has as its primary focus the fast characterization of the progression of sequence data using an interactive graphics system with several controlling parameters.

Path hierarchies in interconnection networks *by P. A. Franaszek, p. 120.* This paper treats the problem of latency minimization in an interconnection network for a system of N high-performance devices. The networks considered here have data transport separated from control, with the data subnetwork designed so that each network function requires only a single control message, and thus only one contention-resolution delay. For sufficiently large N it is shown that (for an abstract hardware model) minimizing contention delays requires that each message subject to such delays have more than one way of reaching its destination (e.g., via a path hierarchy). The overall approach is discussed in the context of the processor-memory interconnection problem in parallel computing.

Volume 31, Number 2, 1987

Preface *by S. Winograd, p. 150.*

Science in industry *by R. E. Gomory, p. 151.* The industrial environment contains problems of practical importance. To exploit that environment, an industrial research group must accept the responsibility for forming strong and useful relationships with the rest of the company of which it is a part. Opportunities for contributing to both the industrial and academic worlds stem from the many product-related problems for which existing science offers no solution, but which suggest new scientific directions.

Reflections on the early days of the department *by H. H. Goldstine, p. 154.* The paper contains the reminiscences of the founding director of the Mathematical Sciences Department in IBM Research. It is interspersed with a number of anecdotes about his friend and colleague, John von Neumann. It attempts in a short space to recapitulate the history of a distinguished department of mathematics in an industrial environment in an informal and hopefully humorous manner.

Applied mathematics, a national view *by H. Cohen, p. 158.* The development of applied mathematics in the United States during the past forty years is described, including the role of computers and computation in enlarging and changing the field. The growth of the profession is also discussed, along with needs for training and funding for the future.

A view of Approximation Theory *by T. J. Rivlin, p. 162.* This is a selective survey of Approximation Theory which touches on the concepts of best approximation, good approximation, approximation of classes of functions, approximation of functionals, and optimal estimation.

Numerical analysis and the scientific method *by J. Glimm and D. H. Sharp, p. 169.* The computer has given rise to a new mode of scientific practice, and today computational science stands beside theory and experiment as a fundamental methodology. The impact of the computer revolution on science can be projected from current trends. The demands to be made on computing methodologies will be reviewed. One of the demands is an ongoing need for excellence in computational methodologies. Generic difficulties encountered in meeting these challenges will be discussed. Recent work of the authors and others will be reviewed in this context.

Some stability techniques for multistep methods *by F. Odeh, p. 178.* A partial survey is given of sundry results in the stability theory of multistep formulae when these are used to integrate time-varying or nonlinear stiff systems, with emphasis on systems arising in circuit analysis.

Factoring logic functions *by R. K. Brayton, p. 187.* A factored form is a representation of a logic function that is either a single literal or a sum or product of factored forms. Thus it is equivalent to a parenthesized algebraic expression. It is one of many possible representations of a logic function, but seems to be the most appropriate one for use in multilevel logic synthesis. We give a number of methods for obtaining different factored forms for a given logic function. These methods range from purely algebraic ones, which are quite fast, to so-called Boolean ones, which are slower but are capable of giving better results. One of the methods given is both fast and gives good results, and is useful in providing continuous estimates of area and delay as logic synthesis proceeds. In multilevel logic synthesis, each of the methods given has a use in a system where run-time and quality are traded off. We also formulate the problem of optimal algebraic factorization, and pose its solution as a rectangle-covering problem for which a heuristic method is given.

Some aspects of the theory of statistical control schemes *by E. Yashchin, p. 199.* Control schemes (charts) are widely used in industrial quality control as means of monitoring the quality of manufactured products. These schemes provide a set of criteria for testing whether a given sequence of observations corresponds to an "on-target" state of the production process. In the present work, we consider some graphical, computational, and statistical aspects of control charting—criteria of performance, methods of derivation, analysis, design, etc. We introduce the class of "Markov-type" control schemes and discuss some of its properties.

The group problem and integer programming duality *by E. L. Johnson, p. 206.* Some duality results for integer programming based on subadditive functions are presented first for linear programs and then for the group problem. A similar result for the knapsack problem is given and then a relationship between facets for the group relaxation and facets of the knapsack problem is given. The mixed integer cyclic group problem is then considered and a dual problem given. A common theme is to try to characterize the strongest possible dual problem or equivalently the smallest possible cone of subadditive functions.

Fourier transforms that respect crystallographic symmetries *by L. Auslander and M. Shenefelt, p. 213.* In crystallography one has to compute finite Fourier transforms that are often very large and often respect crystallographic symmetries. In this paper we discuss efficient finite Fourier transform algorithms on $5 \times 5 \times 5$ points that respect a collection of crystallographic symmetries. Although the size is too small for any practical problems, the methods indicated in this paper can be extended to problems of meaningful size.

The torus and the disk *by R. L. Adler, p. 224.* This paper is a survey of a coherent program of mathematics spanning 28 years. It begins with questions concerning classification and structure in ergodic theory and abstract dynamical systems and describes the author's involvement with toral automorphisms, topological entropy, iteration of maps on the interval, symbolic dynamics, and ultimate engineering applications. It serves as a case study of how unplanned-for practical applications can result from the pursuit of mathematics for its own sake.

The complexity of computations by networks *by N. Pippenger, p. 235.* We survey the current state of knowledge concerning the computation of Boolean functions by networks, with particular emphasis on the addition and multiplication of binary numbers.

Cryptography *by D. Coppersmith, p. 244.* This paper is concerned with two aspects of cryptography in which the author has been working. One is the Data Encryption Standard (DES), developed at IBM and now in wide use for commercial cryptographic applications. This is a "private key" system; the communicants share a secret key, and the eavesdropper will succeed if he can guess this key among its quadrillions of possibilities. The other is the Diffie-Hellman key exchange protocol, a typical "public key" cryptographic system. Its security is based on the difficulty of taking "discrete logarithms" (reversing the process of exponentiation in a finite field). We describe the system and some analytic attacks against it.

Efficient randomized pattern-matching algorithms *by R. M. Karp and M. O. Rabin, p. 249.* We present randomized algorithms to solve the following string-matching problem and

some of its generalizations: Given a string X of length n (the *pattern*) and a string Y (the *text*), find the first occurrence of X as a consecutive block within Y. The algorithms represent strings of length n by much shorter strings called *fingerprints*, and achieve their efficiency by manipulating fingerprints instead of longer strings. The algorithms require a constant number of storage locations, and essentially run in real time. They are conceptually simple and easy to implement. The method readily generalizes to higher-dimensional pattern-matching problems.

Volume 31, Number 3, 1987

Preface *by M. A. Wesley, p. 276.*

Solid modeling for production design *by R. N. Wolfe, M. A. Wesley, J. C. Kyle, Jr., F. Gracer, and W. J. Fitzgerald, p. 277.* This paper describes a solid modeling and interactive graphics computer system which is being used for conceptual and detailed design by the mechanical design community at IBM's Data Systems Division Laboratory in Poughkeepsie, New York. The system has evolved from research on solid modeling begun at IBM's Thomas J. Watson Research Center in the mid-70s. Its development has resulted in one of the first major production uses of solid modeling in industry. The system was first tested in pilot and limited production environments in 1981, and is now in production use as the primary design tool for mechanical portions of IBM's large computer mainframes. Its introduction, development, integration, and use are described, and its functional and performance characteristics as well as requirements for future enhancements are discussed. We conclude from our experience that solid modeling has become a significant new production tool for mechanical design.

Piecewise-circular curves for geometric modeling *by J. R. Rossignac and A. A. G. Requicha, p. 296.* Modern solid modelers must be able to represent a wide class of objects, and must support Boolean operations on solids. These operations are very useful for defining solids, detecting interferences, and modeling fabrication processes. Computing the boundaries of solids defined through Boolean operations requires algorithms for surface/surface and curve/surface intersection. Many of the currently available modelers use closed-form parametric expressions for the curves of intersection of quadric surfaces, and compute intersections of these curves with other surfaces by finding the roots of low-degree polynomials. Because the curves that result from intersections involving tori or more complex surfaces generally cannot be expressed in closed form, modelers typically approximate these curves by cubic splines that interpolate points lying on the true intersections. Cubic splines exhibit second-degree continuity, but they are expensive to process in solid modeling computations. In this paper, we trade second-degree continuity for computational simplicity, and present a method for interpolating three-dimensional points and associated unit tangent vectors by smooth space curves composed of straight line segments and circular arcs. These curves are designated as PCCs (for piecewise-circular curves) and have continuous unit tangents. PCCs can be used in efficient algorithms for performing fundamental geometric computations, such as the evaluation of the minimal distance from a point to a curve or the intersection of a curve and a surface. Formulae and algorithms are presented for generating and processing PCCs in solid modelers. We also show that PCCs are useful for incorporating toroidal primitives, as well as sweeping, growing, shrinking, and blending operations in systems that model solids bounded by the natural quadric surfaces—planes, cylinders, cones, and spheres.

Trimmed-surface algorithms for the evaluation and interrogation of solid boundary representations *by R. T. Farouki, p. 314.* Although trimmed surfaces play a fundamental role in the derivation and processing of solid boundary representations, they have received little attention to date. We propose a trimmed-surface formulation appropriate to the Boolean combination of primitives bounded by a family of elementary surface patches (e.g., planes, quadrics, ruled surfaces, surfaces of revolution) with dual parametric rational polynomial and implicit algebraic equations. Partial intersections between pairs of primitive surface patches are formulated precisely as algebraic curves in the parameter space of each patch. These curves are dissected into monotonic branches by the identification of a characteristic point set. The consolidation of all partial intersections yields a system of piecewise-algebraic loops which define a trimming boundary enclosing a parametric domain for the trimmed patch. With few exceptions, the trimmed-surface formulation is based on precisely defined mathematical procedures, in order to achieve maximum robustness. Some basic interrogation algorithms for solids bounded by trimmed-surface elements are also presented, including procedures for ray-tracing, point/solid classification, sectioning, and computation of surface area, volume, center of gravity, moments of inertia, and other mass properties.

Simple unit vectors orthogonal to a given vector *by M. A. O'Connor and G. Gentili, p. 335.* In geometrical computations it is often necessary to find two unit vectors such that they and a given vector form an orthogonal basis. Computationally simple forms for the two unit vectors are clearly useful. We show that they cannot always be chosen to have rational coordinates, but that in general the simplest possible vectors can be chosen to involve only one square root. We develop number-theoretic criteria for the existence of a rational vector and an effective algorithm for calculation of one if it exists. We also discuss storage and time requirements of the algorithm.

Shaping geometric objects by cumulative translational sweeps *by R. C. Evans, G. Koppelman, and V. T. Rajan, p. 343.* This paper introduces the cumulative translational sweep (CTS) as a tool for shaping geometric objects. It describes

how it may be applied, in combination with Boolean operations, to stimulate growth and shrinking over the boundary regions of polyhedral models, and how, by creating additional facets, it may be used to achieve global rounding effects along model edges and around their vertices. CTSs are examined in terms of a conceptual framework that describes their effects as Minkowski sums—of the polyhedra to be swept, with convex polyhedra from the class of mathematical objects known as zonotopes. Included is a discussion of applications in the OYSTER program, a CAD system for the simulation of semiconductor wafer fabrication.

Voronoi diagram for multiply-connected polygonal domains I: Algorithm *by V. Srinivasan and L. R. Nackman, p. 361.* Voronoi diagrams of multiply-connected polygonal domains (polygons with holes) can be of use in computer-aided design. We describe a simple algorithm that computes such Voronoi diagrams in $O(N(\log_2 N + H))$ time, where N is the number of edges and H is the number of holes.

Voronoi diagram for multiply-connected polygonal domains II: Implementation and application *by S. N. Meshkat and C. M. Sakkas, p. 373.* Voronoi diagrams have many novel applications in computer-aided design. In this paper, an implementation of a Voronoi diagram algorithm described in a companion paper by Srinivasan and Nackman is presented. This Voronoi diagram is then used for an application in which equivalent resistance networks are derived from a boundary representation of a two-dimensional VLSI geometry.

An algorithm for automatic identification of R-fields in bond graphs *by S. J. Hood, R. C. Rosenberg, D. H. Withers, and T. Zhou, p. 382.* Bond graphs may be used to model the power flow in dynamic systems. They are especially attractive for modeling systems which function in coupled energy domains, for example, electromechanical systems. For such systems, bond graphs can be used to provide a natural subdivision into power/energy fields: storage, sources, transformers, and dissipation. In the case of nonlinear dissipative fields, implicit, nonlinear, coupled systems of algebraic equations may arise. Causality assignment on the bond graph provides a basis for detecting implicit formulations. This paper presents an algorithm for detection and solution of these forms within a model, thereby providing an opportunity for efficient numerical solution, and includes a brief introduction to bond graphs via an electromechanical system example.

Parallel algorithms for chip placement by simulated annealing *by F. Darema, S. Kirkpatrick, and V. A. Norton, p. 391.* We explore modifications to the standard simulated annealing method for circuit placement which make it more suitable for use on a shared-memory parallel computer. By employing chaotic approaches we allow the parallel algorithms to deviate from the algorithm defined for a serial computer and thus obtain good execution efficiencies for large numbers of processors. The qualitative behavior of the parallel algorithms is comparable to that of the serial algorithm.

Volume 31, Number 4, 1987

Reducing execution parameters through correspondence in computer architecture *by S. P. Wakefield and M. J. Flynn, p. 420.* The purpose of this study is to develop and extend techniques to provide architectural correspondence between high-level language objects and hardware resources so as to minimize execution time parameters (memory traffic, program size, etc.). A resulting computer instruction set called *Adept* has been emulated, and a compiler has been developed with it as the target language. Although the study was restricted to Pascal, the resulting data are generally applicable to the execution time environment of any procedure-based language. Data indicate that significant bandwidth reductions are possible compared to System/370, VAX, P-code, etc. Specifically, the average improvement ratios realized were instruction bandwidth reduction: 3.46; data read reduction (in bytes): 5.42; data write reduction (in bytes): 14.72.

Performance analysis of the FFT algorithm on a shared-memory parallel architecture *by Z. Cvetanovic, p. 435.* This paper presents a model for the performance prediction of FFT algorithms executed on a shared-memory parallel computer consisting of N processors an the same number of memory modules. The model applies a deterministic analysis to estimate the communication delay through the interconnection network by assuming that all requests arrive at the network in bursts. Our results indicate that the communication delay is significantly affected by the method applied to allocate data to memory modules. For the case in which all data items referenced by a processor during an iteration are allocated to a single memory module, the best-case communication time complexity grows as $O[(\log N)^2/N]$. The worst-case communication time complexity for this case, obtained by a different allocation of data to memory modules, is increased to $O[(\log N)/\sqrt{N}\,]$ due to high network contention. For the case in which the data items referenced by different processors during an iteration are allocated to the same memory module, the communication time complexity is further increased to $O(\log N)$ since all N requests generated by processors are serialized at a single memory module. The methods developed in this paper can be applied for the performance prediction of other well-structured parallel iterative algorithms.

Best and worst mappings for the omega network *by Z. Cvetanovic, p. 452.* This paper presents a study of the best and worst mappings for the omega network proposed by D. H. Lawrie in 1975. We identify mappings that produce no conflicts in the network and mappings that produce a maximum number of conflicts. The analysis of mappings for some typical applications shows that an initial allocation of data to memory modules determines the contention within the

network for all iterations of the algorithm. For the case of the FFT and the bitonic sort algorithm executed on a shared-memory architecture, we prove that if no conflicts are produced during the first iteration of the algorithm, then no conflicts are produced during any other iteration. Moreover, if a maximum number of conflicts are produced during the first iteration, then a maximum number of conflicts are produced during all other iterations of the algorithm. For the d-dimensional grid computations where communication is required with $2d$ nearest neighbors, we prove that if the initial allocation produces no conflicts within the network, then communication with all the neighbors is conflict-free. Also, if the initial allocation produces a maximum number of conflicts, then communication with all the neighbors is maximum-conflict. We show that the omega network cannot produce without conflicts some of the bit-permute mappings such as the perfect shuffle and the bit reversal. The network can produce both of these mappings provided that data items are accessed from memories according to a specific skewed scheme.

Efficient search techniques—An empirical study of the N-Queens Problem *by H. S. Stone and J. M. Stone, p. 464.* This paper investigates the cost of finding the first solution to the N-Queens Problem using various backtrack search strategies. Among the empirical results obtained are the following: 1) To find the first solution to the N-Queens Problem using lexicographic backtracking requires a time that grows exponentially with increasing values of N. 2) For most even values of $N < 30$, search time can be reduced by a factor from 2 to 70 by searching lexicographically for a solution to the $N + 1$-Queens Problem. 3) By reordering the search so that the queen placed next is the queen with the fewest possible moves to make, it is possible to find solutions very quickly for all $N < 97$, improving running time by dozens of orders of magnitude over lexicographic backtrack search. To estimate the improvement, we present an algorithm that is a variant of algorithms of Knuth and Purdom for estimating the size of the unvisited portion of a tree from the statistics of the visited portion.

The inconsistency index method for estimating the accuracy of Schweitzer's approximation *by D. Sitaram, p. 475.* Product-form queueing networks have proved to be useful for predicting the performance of computer systems. In practice, these networks are analyzed using approximate methods because exact methods are computationally too expensive. Schweitzer's approximation is one of the most commonly used. However, there is no method for estimating the error in the solution. This paper proposes a new approach for estimating the error in Schweitzer's approximation for fixed-rate product-form networks. It is based on detecting the extent to which the approximation assumptions used hold. Empirical evidence is presented to show that this approach can be used to accurately predict the error in the approximation.

Exact analysis of round-robin scheduling of services *by H. Takagi, p. 484.* A multi-queue, cyclic-service model with a single-message buffer is considered. Each message consists of a number of characters, and one character is served at the server's each visit. Exact and explicit expressions are derived for performance measures, such as mean cycle time, mean message response time, and mean response time conditioned on the message length. The same model was previously solved by approximation by Wu and Chen [*IBM J. Res. Develop.* 19, No. 5, 486-493 (September 1975)].

Signal degradation through module pins in VLSI packaging *by C.-C. Huang and L. L.-H. Wu, p. 489.* This paper investigates chip-to-chip communication through the modules and board in VLSI packaging. Transmission line models and frequency-dependent transmission line parameters are used in finding the frequency response. The time-domain solution is then obtained through the inverse Fast Fourier Transform. The results show that uncoated module pins, even of relatively short length, can cause severe signal degradation because of their magnetic property. The signal behavior is improved dramatically, however, when the module pins are coated with nonmagnetic conductive material.

Volume 31, Number 5, 1987

Direct semantics of concurrent languages in the SMoLCS approach *by E. Astesiano and G. Reggio, p. 512.* For years providing syntax-directed methods for the formal definition of concurrent languages has proved to be a challenging task. Problems are even more difficult if a language has some of the typical Ada features, such as strong interference between sequential and concurrent aspects, parameterized semantics, complex data structure, and finally an extremely large size. We have developed an approach, the SMoLCS approach, which extends the denotational method to handle concurrent languages and also provides a solution to the above problems. Indeed, our method has been adopted for the format definition of full Ada within the related EEC project. Here we illustrate the basic principles of the approach, following the so-called direct semantics style used for Ada with the help of a toy language as a running example.

Transaction processing primitives and CSP *by J. C. P. Woodcock, p. 535.* Several primitives for transaction processing systems are developed using the notations of Communicating Sequential Processes. The approach taken is to capture each requirement separately, in the simplest possible context: The specification is then the conjunction of all these requirements. As each is developed as a predicate over traces of the observable events in the system, it is also implemented as a simple communicating process; the implementation of the entire system is then merely the parallel composition of these processes. The laws of CSP are then used to transform the system to achieve the required degree of concurrency, to make it suitable for execution in a multiple-tasking system, for example. Finally, there is a

discussion of how state-based systems may be developed using this approach together with some appropriate notation for specifying and refining data structures and operations upon them and of how the system may be implemented. This work is intended as a case study in the use of CSP.

Specification statements and refinement *by C. Morgan and K. Robinson, p. 546.* We discuss the development of executable programs from state-based specifications written in the language of first-order predicate calculus. Notable examples of such specifications are those written using the techniques Z and VDM; but our interest is in the rigorous derivation of the algorithms from which they deliberately abstract. This is, of course, the role of a *development method.* Here we propose a development method based on *specification statements* with which specifications are embedded in programs—standing in for developments "yet to be done." We show that specification statements allow description, development, and execution to be carried out within a *single* language: programs/specifications become hybrid constructions in which both predicates and directly executable operations can appear. The use of a single language—embracing both high- and low-level constructs—has a very considerable influence on the development style, and it is that influence we discuss: the specification statement is described, its associated calculus of refinement is given, and the use of that calculus is illustrated.

A CCS semantics for NIL *by S. A. Smolka and R. E. Strom, p. 556.* We present a syntax-directed translation of NIL, a high-level language for distributed systems programming, into CCS, Milner's Calculus of Communicating Systems. This translation presents unique problems because of NIL's highly dynamic nature, and makes full use of CCS's descriptive facilities. In particular, we consider NIL constructs for dynamic creation and deletion of processes and communication channels, queued synchronous and asynchronous message passing, nondeterministic message selection, and exception handling. A NIL implementation of a simple command shell is used to illustrate the translation procedure. We discuss various issues and open problems concerning the suitability of CCS as an abstract semantics for NIL.

System Identification: An experimental verification *by R. D. Ciskowski, C. H. Liu, H. H. Ottesen, and S.-U.-Rahman, p. 571.* System identification may be defined as the process of determining a model of a dynamic system using observed system input-output data. The identification of dynamic systems through the use of experimental data is of considerable importance in engineering since it provides information about system parameters which is useful in predicting behavior and evaluating performance. Traditional methods of System Identification are usually time-consuming, costly, and difficult to use in other than a product development environment. Within the last decade, more sophisticated techniques for System Identification have been developed that can simultaneously estimate many parameters accurately and repeatedly. These modern techniques are, in addition, efficient, easy to use, inexpensive, and readily adaptable to manufacturing and in-the-field environments where they can be used to evaluate product quality and performance. This paper describes the use of one such System Identification algorithm to estimate several mechanical parameters of 8-inch hard-disk drive spindles in a manufacturing-like setting. The results obtained are in excellent agreement with results acquired by more conventional methods, and demonstrate the potential benefits of System Identification techniques in evaluating product quality and performance.

Delamination and fracture of thin films *by E. Klokholm, p. 585.* The fracture and delamination of thin films is a relatively common occurrence, and prevention of these mechanical failures is essential for the successful manufacture of thin-film devices. Internal elastic stresses are an inherent part of the thin-film deposition process, and are largely responsible for the mechanical failures of thin films. However, it is not the magnitude of the film stress S which governs film fracture or delamination, but the elastic energy U stored in the film. It is the intent of this presentation to show that the mechanical stability of the film and the substrate requires that U be less than a critical value U_c and that U_c is dependent upon the surface energy γ.

Volume 31, Number 6, 1987

Contact metallurgy development for VLSI logic *by R. M. Geffken, J. G. Ryan, and G. J. Slusser, p. 608.* The criteria involved with the choice of an ohmic contact material for VLSI logic are discussed. The problems of aluminum penetration encountered with Al metallization and solid-phase epitaxy associated with Al–Si metallization make these interconnect materials incompatible with VLSI technology. The contact resistance characteristics of palladium and platinum silicides were compared to the contact resistance obtained using a titanium contact layer. The contact resistance of palladium silicide increased with extended annealing at 400°C, while the PtSi and Ti contact materials exhibited stable contact resistance under these conditions. A Ti/Al–Cu/Si process which is compatible with a lift-off patterning technique and partial coverage of contacts is described. Rutherford backscattering results indicate that copper and silicon additions to the aluminum metallization retard the Ti–Al reaction. SIMS data show that silicon in Ti/Al–Cu/Si films redistributes during heat treatment, accumulating at the Ti/Al–Cu interface.

Electrical and microstructural investigation of polysilicon emitter contacts for high-performance bipolar VLSI *by J. M. C. Stork, E. Ganin, J. D. Cressler, G. L. Patton, and G. A. Sai-Halasz, p. 617.* Key electrical characteristics of polysilicon emitter contacts in bipolar transistors, such as contact resistance and recombination velocity, are extremely sensitive to the microstructure of the polysilicon/single-crystal

silicon interface. In this study, we correlated the microstructural and electrical characteristics of this interface by performing cross-sectional transmission electron microscopy (XTEM) on *actual* transistors on the same chip where ring-oscillator speeds were measured. The base current and emitter resistance of the fastest devices approached values typical of single-crystal silicon emitters. Interpretation of these electrical data and of the SIMS impurity profile indicates that significant restructuring of the polysilicon/single-crystal interface had taken place. This conclusion was indeed confirmed by the XTEM results. Although the low-current performance was degraded because of higher junction capacitances, the high-current switching speed was improved because of the minimal emitter contact resistance. Since the current gain was sufficiently high and very uniform, it is concluded from this work that minimization of both junction depth and contact resistance is the most important design consideration for high-performance submicron transistors, rather than maximization of the gain enhancement of the polysilicon/single-crystal interface.

Study of contact and shallow junction characteristics in submicron CMOS with self-aligned titanium silicide *by Y. Taur, B. Davari, D. Moy, J. Y.-C. Sun, and C.-Y. Ting, p. 627.* The contact resistance between $TiSi_2$ and shallow n^+/p^+ source-drains in CMOS is studied for a variety of junction depths and silicide thicknesses. The contact contribution to the total device series resistance can be significant if excessive silicon and dopants are consumed during silicide formation. Low contact resistances are obtained for 0.15-μm n^+ and 0.20-μm p^+ junctions when the titanium thickness is reduced to keep a high doping concentration at the $TiSi_2$/Si interface. Alternatively, a nonstandard process can be employed to implant additional dopants into the titanium. A thin layer of dopants then out-diffuses into the silicon after the silicide reaction and anneal to help reduce contact resistance and leakage currents. The latter technique is more extendable to CMOS devices which require thicker titanium films and/or shallower junctions.

Oxidation of Si-rich chemical-vapor-deposited films of tungsten silicide *by L. Krusin-Elbaum and R. V. Joshi, p. 634.* We have studied dry oxidation characteristics of Si-rich WSi_x thin films prepared by LPCVD directly on SiO_2, with $x = 2.7$ for as-deposited films. It has been reported previously that thin (less than 100 nm) CVD tungsten silicide adheres well to SiO_2. Using Auger depth profiling and Rutherford backscattering spectroscopies, we find that silicon in excess of stoichiometric WSi_2 diffuses through the silicide toward the surface to form a SiO_2 passivating overlayer. The extracted activation energy for this oxidation process is $E_a = 1.2$ eV, consistent with oxygen diffusion in SiO_2. A similar value of E_a is found for WSi_x deposited on polysilicon. During the anneal, the stoichiometry x of WSi_x decreases monotonically with the annealing temperature, reaching $x = 2$ after 30 min at 900°C or 20 min at 950°C. Longer times or higher temperatures result in silicon depletion, with $x = 1.7$ after 30 min at 1000°C. At the same time, the resistivity of WSi_x also decreases from ≈90 Ω/□ for 1500 Å as-deposited film to 5 Ω/□ for the films annealed at 1000°C, the value obtained in a standard homogenization anneal. A scanning electron micrograph (SEM) of 0.5-μm fine lines patterned using e-beam lithography reveals that the integrity of fine line structures, their adhesion to SiO_2, and their vertical profiles remain unchanged after the oxidation process. We suggest that such Si-rich tungsten silicide can be useful as a gate electrode without the polysilicon underlayer, since no extra passivation is necessary and reoxidation and homogenization steps in the FET processing sequence can be accomplished simultaneously.

Correlation analysis of particle clusters on integrated circuit wafers *by C. H. Stapper, p. 641.* Defect clustering results in correlations between the numbers of defects or faults that occur on integrated circuit chips located adjacent to one another on semiconductor wafers. Until now, it has been believed that correlations of this type were not accounted for in existing yield models. It is shown in this paper that such correlations are present in yield models based on mixed or compound Poisson statistics. A quadrat analysis of particle distributions on semiconductor wafers is used to compare data and theory. The results show that the theoretical correlation coefficients are in agreement with the experimental ones. It was also determined from the particle data how these correlation coefficients vary as the distance between quadrats is increased. This variation provides a convenient method for determining the cluster dimensions.

Modeling and analysis of computer system availability *by A. Goyal and S. S. Lavenberg, p. 651.* The quantitative evaluation of computer-system availability is becoming increasingly important in the design and configuration of commercial computer systems. This paper deals with methods for constructing and solving large Markov-chain models of computer-system availability. A set of powerful high-level modeling constructs is discussed that can be used to represent the failure and repair behavior of the components that comprise a system, including important component interactions, and the repair actions that are taken when components fail. If time-independent failure and repair rates are assumed, then a time-homogeneous continuous-time Markov chain can be constructed automatically from the modeling constructs used to describe the system. Markov chains having tens of thousands of states can be readily constructed in this manner. Therefore, techniques that are particularly suitable for numerically solving such large Markov chains are also discussed, including techniques for computing the sensitivities of availability measures with respect to model parameters. A computer system modeling example is presented to illustrate the use of these modeling and analysis techniques. The modeling constructs, automatic Markov-chain construction, and model-solution methods have been implemented in a program package called the System Availability Estimator (SAVE).

Ferroresonance *by D. M. Scoggin and J. E. Hall, Jr., p. 665.* This paper describes a mathematical model for ferroresonant circuits that addresses some of the deficiencies of earlier analyses of ferroresonant regulators. Derived using piecewise-linear, normalized differential equations, the model accommodates nonlinear behavior and predicts circuit performance in terms of parameters such as line voltage, frequency, and load. A phase-plane analysis is used to simplify the determination of linear regions of operation between nonlinear events. Numerical solutions of the resulting equations are used to generate time-domain and parametric performance curves. The results compare well with experiments and suggest potential applications in the design of high-frequency voltage regulators.

Font design for high-speed impact line printers *by J. L. Zable and H. C. Lee, p. 679.* In impact line printers that use print-band (or similar) technology, the higher speed required of the type band for higher print throughput results in wider printed strokes with increased slur. Ordinarily, font designers compensate for the increased printed strokewidth by narrowing the width of the engraved characters on the type band. While this approach corrects the total printed character stroke, the print quality is degraded because of increased slur. This paper presents an alternative design approach in which an examination of the essential parameters of print dynamics suggests a font design that incorporates wider strokewidths.

Volume 32, Number 1, 1988

Preface *by M. Büttiker, p. 3.*

World as system self-synthesized by quantum networking *by J. A. Wheeler, p. 4.* The quantum, strangest feature of this strange universe, cracks the armor that conceals the secret of existence. In contrast to the view that the universe is a machine governed by some magic equation, we explore here the view that the world is a self-synthesizing system of existences, built on observer-participancy via a network of elementary quantum phenomena. The elementary quantum phenomenon in the sense of Bohr, the elementary act of observer-participancy, develops definiteness out of indeterminism, secures a communicable reply in response to a well-defined question. The rate of carrying out such yes-no determinations, and their accumulated number, are both minuscule today when compared to the rate and number to be anticipated in the billions of years yet to come. The coming explosion of life opens the door, however, to an all-encompassing role for observer-participancy: to build, in time to come, no minor part of what we call *its* past—*our* past, present, and future—but this whole vast world.

Notes on the history of reversible computation *by C. H. Bennett, p. 16.* We review the history of the thermodynamics of information processing, beginning with the paradox of Maxwell's demon; continuing through the efforts of Szilard, Brillouin, and others to demonstrate a thermodynamic cost of information acquisition; the discovery by Landauer of the thermodynamic cost of information destruction; the development of the theory of and classical models for reversible computation; and ending with a brief survey of recent work on quantum reversible computation.

Miniaturization of electronics and its limits *by R. W. Keyes, p. 24.* The long-continued advance of the performance of information processing technologies has been based on miniaturization of components. The history of miniaturization is presented through examples. They suggest that limits proposed by Landauer in the 1960s will be reached in two or three decades.

Information transport obeying the continuity equation *by T. Toffoli, p. 29.* We analyze nontrivial dynamical systems in which information flows as an additive conserved quantity—and thus takes on a strikingly tangible aspect. To arrive at this result, we first give an explicit characterization of equilibria for a family of lattice gases.

Origin of life and physics: Diversified microstructure—Inducement to form Information-Carrying and Knowledge-Accumulating systems *by H. Kuhn, p. 37.* The process leading to the origin and evolution of life is caused by the presence of distinct physical and chemical conditions at a distinct location in the universe. A specified system originates and evolves under the continuous influence of a complex operational environment. The system develops toward increasing independence of the original environment by becoming increasingly complex. Modeling a detailed scenario consisting of a sequence of reasonable physico-chemical steps is essential in rationalizing the phenomenon. The basic process, accumulation of knowledge by continuously testing environmental properties, is intimately related to the measuring process in physics. Evolution is a physical process, and this process leads to man developing physics. Thus physics appears to be self-consistent—the basis and consequence of evolution. The physics-producing system is considered to be a measuring and information-processing device based upon the mechanism which operates in the origin and evolution of life.

Density of states of one-dimensional random potentials *by P. Erdös, p. 47.* Following an introduction to the early history of the theories of the density of electronic states in one-dimensional structures, pioneered, among others, by R. Landauer and J. C. Helland, a particular model, that of a multistep random potential, is discussed. It is shown that Kolmogorov-type equations can be obtained for the probability distribution of the phase of the wave function, and, by solving these equations, the density of states may be calculated. An analogy with the classical rotator in a random force field is worked out, and helps in visualizing the results.

Bloch electron in a magnetic field: Mixed dimensionality and the magnetic-field-induced generalized quantum Hall effect *by M. Y. Azbel, p. 52.* The energy spectrum of a Bloch electron in a magnetic field is one-dimensional. This leads to the Peierls instability and the magnetic-field-induced transition to the quantized Hall effect. The wave function is two-dimensional. This decreases the Peierls gap and makes it exponentially vanishing with magnetic field. Disorder lifts the degeneracy and one-dimensionality of the spectrum. High disorder yields a metallic behavior. Intermediate disorder leads to the generalized quantized Hall effect. The latter has a finite magnetoresistance as a semimetal, and Hall plateaus similar to the quantized ones, but they may have any value of the effective charge.

Residual resistivity dipoles, electromigration, and electronic conduction in metallic microstructures *by R. S. Sorbello and C. S. Chu, p. 58.* For an impurity in a bulk metal, the connection between electromigration and electric fields associated with dc conductivity is understood in terms of Landauer's residual resistivity dipole. This connection is examined, and appropriate generalizations are made for an impurity in a two-dimensional electron gas and for an impurity near a metal surface. The residual resistivity dipole field decays less rapidly with distance in a two-dimensional gas than in bulk, thus resulting in a larger voltage drop across an impurity in the system of lower dimensionality.

Coherent and sequential tunneling in series barriers *by M. Büttiker, p. 63.* A simple approach which can describe

both coherent tunneling and sequential tunneling is applied to resonant tunneling through a double-barrier structure. This approach models phase-randomizing events by connecting to the conductor a side branch leading away from the conductor to a reservoir. The reservoir does not draw or supply a net current, but permits inelastic events and phase randomization. A conductance formula is obtained which contains contributions due to both coherent and sequential tunneling. We discuss the limiting regimes of completely coherent tunneling and completely incoherent transmission, and discuss the continuous transition between the two. Over a wide range of inelastic scattering times tunneling is sequential. The effect of inelastic events on the peak-to-valley ratio and the density of states in the resonant well is investigated. We also present an analytic discussion of the maximum peak conductance e/h of an isolated resonance in a many-channel conductor.

Diffusion from an entrance to an exit *by M. E. Fisher, p. 76.* Asymptotic and exact solutions are derived from first principles by various methods for the moments of the number of steps or traversal time, etc., of a particle which diffuses, most specifically on a linear chain, to an exit site without previously leaving via an entrance site. The presentation is expository and uses standard methods.

Band tails, path integrals, instantons, polarons, and all that *by M. H. Cohen, M.-Y. Chou, E. N. Economou, S. John, and C. M. Soukoulis, p. 82.* This paper reviews the explanations recently developed by the authors and their collaborators of how disorder leads to exponential band tails and to Urbach tails in optical absorption. It starts with the simplest single-potential-well models which, despite their simplicity, are remarkably successful in accounting for the experimental facts. It then identifies the weaknesses, hidden or explicit, in these models and shows, step by step, how they can be corrected by increasing the sophistication of the procedures used. Exact results are finally achieved through use of field-theoretic techniques, and appropriately formulated single-potential-well models are shown to reproduce these quite accurately. It is also shown that the probability distribution of the random potential must be close to Gaussian, with an autocorrelation function which cuts off fairly rapidly with distance for there to be a well-defined, broad energy range in which there are exponential band tails in the density of states and Urbach tails in the optical absorption.

Fundamental questions in the theory of electromigration *by A. H. Verbruggen, p. 93.* The theory of electromigration is focused on the force acting on a lattice defect in a metallic sample that carries an electric current. Much work has been done to obtain a better understanding of the underlying physical mechanisms. The force has been calculated numerically for defects in several metals, and a qualitative agreement with the experiments has often been found. There are, however, still discussions about the relevance of certain contributions to the force. These originate from conceptual difficulties related to 1) the nature of the screening of the electric field at the site of an impurity by the conduction electrons and 2) the existence and significance of inhomogeneities in the electric field and current flow near an impurity. This paper provides a review of the basic models and of questions which still exist in the theory of electromigration. The relevance of these questions is illustrated by results of experimental work on the electromigration of H in V, Nb, and Ta.

Tunneling times and a quantum clock *by C. Foden and K. W. H. Stevens, p. 99.* The problem of measuring tunneling times by means of a quantum clock is found to lead to difficulties which are thought to arise because the Hamiltonian of the coupled system does not separate into particle and clock parts.

Coherent voltage oscillations in small normal tunnel junctions and the crossover to the incoherent regime *by Y. Gefen, p. 103.* We discuss the possibility of charge oscillations in a normal tunnel junction, driven by an external current source (I_{ex}), in the coherent limit. In that limit the dephasing time t_ϕ is larger than the period $t_p \equiv e/I_{ex}$. This behavior is modified when t_ϕ decreases.

Relative stability in nonuniform temperature *by N. G. van Kampen, p. 107.* Landauer has suggested that the relative stability of a particle diffusing in a bistable potential is affected by an intervening hot layer. We derive this effect both from thermodynamics and from the diffusion equation. For this purpose the proper form of the diffusion equation in a nonuniform medium is established for the case of a Brownian particle. If the diffusion takes place in a ring, the hot layer creates a steady current.

Boundary-layer theory for the extremely underdamped Brownian motion in a metastable potential *by H. Risken, K. Vogel, and H. D. Vollmer, p. 112.* A theory for the boundary layer near the critical trajectory for the extremely underdamped Brownian motion in a metastable potential is presented. The probability distribution function in phase space near this critical trajectory, the average escape energy, and the correction terms for the zero-friction-limit escape rate are calculated.

Bistability in active circuits: Application of a novel Fokker–Planck approach *by R. Hänggi and P. Jung, p. 119.* The problem of metastability in electronic circuits with negative differential resistance, originally pioneered by Landauer in 1962, is reconsidered from the viewpoint of a Fokker–Planck modeling for nonlinear shot noise (master equation). A novel Fokker–Planck approximation scheme is presented that describes correctly the deterministic flow and the long-time dynamics of the master equation. It is demonstrated that the conventional scheme of a truncated Kramers–Moyal expansion at the second order overestimates the transition rates in leading exponential order. In order to obtain the correct relative stability, the novel scheme uses a

diffusion coefficient which incorporates information about global nonlinear fluctuations characterized by the whole set of higher-order Kramers–Moyal transport coefficients.

Quantum noise and quantum Langevin equations *by C. W. Gardiner, p. 127.* The quantum Langevin equation of Ford, Kac, and Mazur is rederived and shown to be equivalent to an adjoint equation. This latter can be handled by means of van Kampen's cumulant expansion to yield derivations of the quasiclassical Langevin equation, stochastic electrodynamics, quantum optical, and quantum Brownian motion master equations (under appropriate conditions). The result of Benguria and Kac—that the quantum Langevin equation yields the Boltzmann distribution over energy levels in thermodynamic equilibrium—is also verified.

Symmetry and transport in disordered systems *by J. Pendry, p. 137.* The transfer matrix for a disordered system enables averages of integer powers of the resistance to be found, R^N; application of the symmetric group generalizes this formula to fractional and negative N, providing a powerful tool for the study of transport. Consequences for fluctuations in resistance, $1/f$ noise, and frequency response are discussed, as well as a new sort of state, of fractal dimension 1/2, which is responsible for transport in localized systems.

Correlated discrete transfer of single electrons in ultrasmall tunnel junctions *by K. K. Likharev, p. 144.* Recent theoretical and experimental studies have revealed a new family of effects taking place in very small tunnel junctions at low temperatures. The effects have a common origin, the correlated discrete tunneling of single electrons and/or Cooper pairs resulting from their electrostatic ("Coulomb") interaction. This paper presents a brief review of the single-electron part of the family, including discussion of the background physics, methods of theoretical description of the new effects, experimental results, and possible applications of the new effects in analog and digital electronics.

Volume 32, Number 2, 1988

A flexible graph-unification formalism and its application to natural-language processing *by G. Bouma, E. König, and H. Uszkoreit, p. 170.* A graph-unification-based representation language is described that was developed as the grammar formalism for the LILOG research project at IBM Germany. The Stuttgart Unification Formalism (STUF) differs from its predecessors in its higher flexibility and its algebraic structure. It is well suited for the implementation of rather different linguistic approaches, but is currently employed mainly in the development of Categorial Unification Grammars with a lexicalized compositional semantics. Examples from the syntactic and semantic processing of natural language are used to illustrate the virtues of the formalism and of our lexicalist approach to linguistic analysis.

An experiment in computational discrimination of English word senses *by E. Black, p. 185.* A number of researchers in text processing have independently observed that people can consistently determine in which of several given senses a word is being used in text, simply by examining the half dozen or so words just before and just after the word in focus. The question arises whether the same task can be accomplished by mechanical means. Experimental results are presented which suggest an affirmative answer to this query. Three separate methods of discriminating English word senses are compared information-theoretically. Findings include a strong indication of the power of domain-specific content analysis of text, as opposed to domain-general approaches.

Spelling assistance for compound words *by R. Frisch and A. Zamora, p. 195.* This paper describes a method for providing spelling assistance for Germanic compound words. The technique systematically analyzes an unknown word to determine its components, using a dictionary which associates word components with codes that describe their compounding characteristics. Language-specific morphological transformations are used to take into consideration common intraword elision patterns. Special dictionary entries, heuristic rules, and lexical distance measures are used to provide the best possible replacement compound words. The method is fast and provides spelling assistance and hyphenation support in an interactive environment.

CRITAC—An experimental system for Japanese text proofreading *by K. Takeda, E. Suzuki, T. Nishino, and T. Fujisaki, p. 201.* This paper describes an experimental expert system for proofreading Japanese text. The system is called CRITAC (CRITiquing using ACcumulated knowledge). It can detect typographical errors, Kana-to-Kanji conversion errors, and stylistic errors in Japanese text. We describe the basic concepts and features of CRITAC, including preprocessing of text, a high-level text model, Prolog-coded heuristic proofreading knowledge, and a user-friendly interface. Although CRITAC has been primarily designed for Japanese text, it appears that most of the concepts and the architecture of CRITAC can be applied to other languages as well.

Large-vocabulary speech recognition: A system for the Italian language *by P. D'Orta, M. Ferretti, A. Martelli, S. Melecrinis, S. Scarci, and G. Volpi, p. 217.* We describe a research project in automatic speech recognition which has led to the development of an experimental large-vocabulary real-time recognizer for Italian, and show how the maximum-likelihood techniques which had been employed in the development of prototype recognizers for English can be tailored to a language with substantially different characteristics.

Multilevel decoding for Very-Large-Size-Dictionary speech recognition *by B. Mérialdo, p. 227.* An important concern in the field of speech recognition is the size of the vocabulary

that a recognition system is able to support. Large vocabularies introduce difficulties involving the amount of computation the system must perform and the number of ambiguities it must resolve. But, for practical applications in general and for dictation tasks in particular, large vocabularies are required, because of the difficulties and inconveniences involved in restricting the speaker to the use of a limited vocabulary. This paper describes a new organization of the recognition process, Multilevel Decoding (MLD), that allows the system to support a Very-Large-Size Dictionary (VLSD)—one comprising over 100,000 words. This significantly surpasses the capacity of previous speech-recognition systems. With MLD, the effect of dictionary size on the accuracy of recognition can be studied. In this paper, recognition experiments using 10,000- and 200,000-word dictionaries are compared. They indicate that recognition using a 200,000-word dictionary is more accurate than recognition using a 10,000-word dictionary (when unrecognized words are included in the error rate).

A Japanese sentence analyzer *by N. Maruyama, M. Morohashi, S. Umeda, and E. Sumita, p. 238.* This paper presents the design of a broad-coverage Japanese sentence analyzer which can be part of various Japanese processing systems. The sentence analyzer comprises two components: the lexical analyzer and the syntactic analyzer. Lexical analysis, i.e., segmenting a sentence into words, is a formidable problem for a language like Japanese, because it has no explicit delimiters (blanks) between written words. In practical applications, this task is made more difficult by the occurrence of words not listed in a dictionary. We have developed a five-layered knowledge source and used it successfully in the lexical analyzer, resulting in very accurate segmentation, even in cases where there are unknown words. The syntactic analyzer has two modules: One consists of an augmented context-free grammar and the PLNLP parser; the other is the dependency structure constructor, which converts the phrase structures to dependency structures. The dependency structures represent various key linguistic relations in a more direct way. The dependency structures have semantically important information such as tense, aspect, and modality, as well as preference scores reflecting relative ranking of parse acceptability.

Conceptual graphs for the analysis and generation of sentences *by P. Velardi, M. T. Pazienza, and M. De'Giovanetti, p. 251.* A system for analyzing and generating Italian texts is under development at the IBM Rome Scientific Center. Detailed semantic knowledge on word-sense patterns is used to relate the linguistic structure of a sentence to a conceptual representation (a *conceptual graph*). Conceptual graphs are stored in a database and accessed by a natural-language query/answering module. The system analyzes a text supplied by a press-agency-release database. It consists of three modules: a morphological, a syntactic, and a semantic processor. The semantic analyzer uses a *conceptual lexicon* of word-sense descriptions, currently including about 850 entries. A description is an extended case frame providing the *surface semantic patterns* (SSP) of a word-sense *w*. SSPs express both semantic constraints and word-usage information, such as commonly found word patterns, idioms, and metaphoric expressions. SSPs are used by the semantic interpreter to build a conceptual graph of the sentence, which is then accessed by the query-answering and language-generation modules. This paper makes the claim that the SSP approach is viable and necessary to cope with language phenomena in unrestricted domains. Surface patterns are easily acquired inductively from the natural-language corpus rather than deductively from predefined conceptual structures. SSPs map quite complex sentences into surface semantic representations that can be generalized at a subsequent stage. In contrast, the current state of the art does not provide viable theory or methodology to go from superficial to deep structures. This issue is more extensively addressed in the body of the paper.

An analysis of hardware and software availability exemplified on the IBM 3725 Communication Controller *by P. I. Pignal, p. 268.* Because of the growing commercial, governmental, and scientific requirements for system availability, evaluating this factor has become increasingly important. This paper presents a unified approach to hardware and software availability of a system in the operational phase. The aim is to evaluate the availability in a given time interval, to show how to improve it, and to determine the probability that a specified level is met over the period. The inputs are the failure and repair rates of the system elements, and the functional relationship between them. Field tracking provides the failure and repair data, and Markov-chain techniques make it possible to construct, reduce, and solve the model. Availability is computed by the program package System Availability Estimator (SAVE). The model has been used and validated with actual field data for the IBM 3725 Communication Controller.

Volume 32, Number 3, 1988

Preface *by R. A. Webb and R. B. Laibowitz, p. 304.*

Spatial variation of currents and fields due to localized scatterers in metallic conduction *by R. Landauer, p. 306.* Volume 1 of this journal, thirty-one years ago, included a paper with the above title. Studies of small samples, in recent years, as well as earlier work on disordered samples, have caused some of the content of the earlier work to become widely understood. The aspects stressed in the title, however, relating to the spatial variations in the vicinity of a localized scattering center, have received little attention, except in electromigration theory debates. Here, we return to these aspects of the earlier paper, and emphasize that the transport field associated with a point-defect scattering center is a highly localized dipole field. The nonlinearity of resistance in terms of scattering cross section is discussed. A theory of these effects, which does justice to the coherent

multiple-scattering effects present at low temperatures, does not yet exist. Such a theory is likely to modify the effects, but it is unlikely to cause them to disappear. We also discuss closed loops, without leads; the persistent currents expected in these; and a possible method of detecting the persistent currents.

Symmetry of electrical conduction *by M. Büttiker, p. 317.* The resistance of a conductor measured in a four-probe setup is invariant if the exchange of the voltage and current sources is accompanied by a magnetic field reversal. We present a derivation of this theorem. The reciprocity of the resistances is linked directly to the microscopic reciprocity of the S-matrix, which describes reflection at the sample and transmission through the sample. We demonstrate that this symmetry holds for a conductor with an arbitrary number of leads. Since leads act like inelastic scatterers, consideration of a many-probe conductor also implies that the reciprocity of resistances is valid in the presence of inelastic scattering. Various conductance formulae are discussed in the light of the reciprocity theorem. Finally, we discuss some implications of our results for the nature of a voltage measurement and point to the difference between chemical potentials and the local electric potential.

Fluctuations in the extrinsic conductivity of disordered metal *by S. Washburn, p. 335.* Random fluctuations of the electrical conductance are ubiquitous in small (typical dimension $L < 1$ μm) metallic samples at low temperatures (typically $T < 1\text{K} \approx 0.09$ meV). The fluctuations result from the quantum-mechanical interference of the carrier wavefunctions. The superpositions of the wavefunctions depend randomly on the placement of impurities, on magnetic field, and on the current driven through the sample. At length scale L_ϕ (the average distance over which the carriers retain phase information), the fluctuations always have amplitude $\Delta G \approx e^2/h$, and any observations at scale larger than the coherence length yield a decreased amplitude of the fluctuations. Since the carrier wavefunctions are not classical, local objects (they extend over regions of size L_ϕ), the conductance contains nonlocal terms. For instance, the conductance is not zero far from the classical current paths through the sample and is not symmetric under the reversal of the magnetic field. In this article, the physics of the fluctuations is reviewed, and some of the experiments which illuminate the physics are described.

Mesoscopic coherence phenomena in semiconductor devices *by S. B. Kaplan and A. Hartstein, p. 347.* Semiconductor devices have several attractive properties which make them useful in the study of electronic coherence phenomena such as universal conductance fluctuations. The use of gated devices allows the Fermi level, and thus the electronic wavelength, to be adjusted in order to study energy correlation effects. The two-dimensional electron gas formed beneath the gate can be tilted with respect to the magnetic field to reveal that the field correlation length of the fluctuations obeys a cosine law. This strongly suggests that the fluctuations are caused by quantum interference in the same way that the Aharonov–Bohm effect arises in metallic rings. The energy range over which electrons are correlated in these materials is generally larger than in metals. This allows one to study these conductance fluctuations at much higher temperatures than are feasible in metallic conductors. For the same reason, substantially larger source–drain voltages can be applied to observe asymmetry and nonlinear effects in the conductance.

Isolated rings of mesoscopic dimensions. Quantum coherence and persistent currents *by H.-F. Cheung, Y. Gefen, and E. K. Riedel, p. 359.* Persistent currents in small nonsuperconducting rings threaded by a magnetic flux are a manifestation of novel quantum effects in submicron systems. We present theoretical results for one-channel and multichannel systems concerning the dependence of the current amplitude on the number of channels and geometry, temperature, and degree of disorder. Inelastic scattering is considered for one-channel loops only. We also discuss the observability of the effect.

Electronic transport in small strongly localized structures *by A. B. Fowler, J. J. Wainer, and R. A. Webb, p. 372.* We review some recent results on the low-temperature transport properties ($T < 4$K) of very small silicon metal-oxide field-effect transistors in the insulating regime of conduction. Our devices are lithographically patterned to have widths as small as 0.05 μm and lengths as short as 0.06 μm. These small transistors exhibit new and unexpected sample-specific fluctuation behavior in the gate voltage, temperature, and magnetic field dependence of the conductance. We discuss both resonant tunneling and Mott variable-range hopping, the two main transport mechanisms in these devices at low temperature.

What is measured when you measure a resistance?—The Landauer formula revisited *by A. D. Stone and A. Szafer, p. 384.* We re-examine the question of what constitutes the physically relevant quantum-mechanical expression for the resistance of a disordered conductor in light of recent experimental and theoretical advances in our understanding of the conducting properties of mesoscopic systems. It is shown that in the absence of a magnetic field, the formula proposed by Büttiker, which expresses the current response of a multi-port conductor in terms of transmission matrices, is derivable straightforwardly from linear response theory. We also present a general formalism for solving these equations for the resistance given the scattering matrix. This Landauer-type formula reduces to $g = (e^2/h)Tr(tt^\dagger)$, where g is the conductance and t is the transmission matrix, for the two-probe case. It is suggested that this formula provides the best description of the present class of experiments performed in two-probe or multi-probe measuring configurations, and that the subtleties leading to various different Landauer formulae are not relevant to these experiments. This is not

because of the large number of channels in real conductors, but is due to the fact that apparently no present experiment probes a "local chemical potential" in the conductor. Certain standard objections to deriving a Landauer-type formula from linear response theory are answered. Applications of this formula to fluctuations in disordered multi-probe conductors are discussed.

Scanning tunneling measurements of potential steps at grain boundaries in the presence of current flow *by J. R. Kirtley, S. Washburn, and M. J. Brady, p. 414.* We have used a new technique to simultaneously measure the surface topography and surface potential of current-carrying polycrystalline $Au_{60}Pd_{40}$ thin films using a scanning tunneling microscope. We find abrupt steps in the surface potential due to scattering from grain boundaries in these films.

Volume 32, Number 4, 1988

Preface *by R. B. Laibowitz and R. A. Webb, p. 440.*

Materials and processes for microstructure fabrication *by M. Hatzakis, p. 441.* The fabrication of structures considerably smaller than the devices and circuits that are mass-produced for use in computers and other electronic equipment is the subject of this paper. Devices of < 1 μm (microstructures) and < 100 nm (nanostructures) minimum dimensions were made possible in a practical sense only after the introduction of electron beams and the associated processes, as lithographic tools in the early 1960s. This paper presents a historical perspective of this very important chapter in lithographic technology, primarily from the point of view of materials and processes, since electron-beam systems are covered in other papers in this issue. In addition, the important criteria that have to be considered in the fabrication of small structures, with respect to the interaction of the writing beam with the resist material and the substrate, and the subsequent pattern-transferring processes, are discussed.

Nanolithography with a high-resolution STEM *by C. P. Umbach, A. N. Broers, R. H. Koch, C. G. Willson, and R. B. Laibowitz, p. 454.* A high-resolution scanning transmission electron microscope (STEM) with a beam diameter approaching 0.6 nm has been adapted for the patterning of complex fine-line nanostructures. An IBM PC XT is used as the pattern generator to direct the scan electronics from a Cambridge Stereoscan 250 which have been interfaced with the scanning coils of the STEM. A study of the ultimate resolution of the newly designed acid-catalyzed resist poly(p-t-butyloxycarbonyloxystyrene) has been carried out. The STEM has proven to be a flexible tool in the fabrication of individual nanostructure devices for quantum transport studies in mesoscopic devices smaller than an electron phase-coherence length.

Nanostructure technology *by T. H. P. Chang, D. P. Kern, E. Kratschmer, K. Y. Lee, H. E. Luhn, M. A. McCord, S. A. Rishton, and Y. Vladimirsky, p. 462.* The ability to fabricate structures with lateral dimensions in the sub-100-nm range has opened a new field of research. This paper first reviews recent advances in nanolithography techniques, with a brief discussion of their relative merits and fundamental limits. Special emphasis is given to the scanning electron-beam method, which is the most widely used nanolithography method at the present time. The two main areas of nanostructure research are device technology and basic science. Highlights of a number of exploratory programs in these two areas are presented.

High-throughput, high-resolution electron-beam lithography *by H. C. Pfeiffer, T. R. Groves, and T. H. Newman, p. 494.* The introduction of the shaped-beam imaging technique has greatly enhanced the exposure efficiency of electron-beam lithography systems. IBM's EL systems provide the throughput needed for lithography applications in semiconductor fabrication lines. The resolution of these systems has been steadily improved over the past 15 years in support of the semiconductor lithography trend toward submicron dimensions. This paper describes the latest version (EL-3 system) capable of fabricating 0.25-μm features. The technical challenges of submicron e-beam lithography are discussed, and practical solutions together with experimental results are presented.

Resolution limits for electron-beam lithography *by A. N. Broers, p. 502.* This paper discusses resolution limits for electron-beam fabrication. Electron beams have been used to produce structures 1 nm in size and useful devices with minimum features of about 20 nm. In all cases the resolution is set primarily by the range of the electron interaction phenomena that form the structures, and not by the size of the electron beam used to write the patterns. The beam can be as small as 0.5 nm. All useful devices built to date have been fabricated with conventional resist processes; these have an ultimate resolution of about 10 nm. Experimental data for PMMA, the highest-resolution electron resist, show that resolution is independent of molecular weight and is therefore not a function of the molecular size. The most promising of the methods offering resolution below 10 nm is the direct sublimation of materials such as AlF_3 and Al_2O_3; 1-nm structures have been fabricated, but it has not been possible to convert the structures into useful devices. In addition to the processes which use intermediate patterned layers, there is the possibility of making devices by direct modification of the electrical properties of conductors, semiconductors, or superconductors by means of high-energy electron bombardment. In these cases no intermediate fabrication process would be used, and it might be possible to reach dimensions comparable to the beam diameter.

Advanced electron-beam lithography for 0.5-μm to 0.25-μm device fabrication *by F. J. Hohn, A. D. Wilson, and P. Coane, p. 514.* High-resolution lithographic capability is required for the fabrication of fully scaled semiconductor

devices at minimum dimensions of 0.5 μm to 0.25 μm—the prototype for the semiconductor logic and memory CMOS devices of the 1990s. Electron-beam exposure tools provide this capability. Fully scaled 0.5-μm test devices were fabricated using a modified EL-3 variable shaped-electron-beam system, while 0.25-μm ground-rule lithography was accomplished with a Gaussian round-electron-beam Vector Scan system. An important part of this technology is the selection of lithographic resist system and the process used for pattern definition and transfer. Twelve or more lithographic steps are often needed for circuit devices with the above minimum dimensions. For fully scaled applications, each one of these pattern levels must be defined by electron-beam lithography, and each level may require a specific lithographic resist. Thus, the electron-beam system and the resist process must be mutually compatible if the required resolution, feature size control, and pattern-level-to-pattern-level overlay accuracy are to be achieved. This paper discusses the successful integration of e-beam lithography and resist technologies and their application to CMOS device fabrication.

Analysis of page-reference strings of an interactive system *by M. G. Kienzle, J. A. Garay, and W. H. Tetzlaff, p. 523.* The performance of real-storage-management algorithms in interactive systems suggests that locality of reference extends to a significant degree across users' transactions. This paper investigates this locality of reference by analyzing page-reference strings gathered from live systems. The data confirm the supposition: They suggest that reference patterns are dominated by system data references that are implied by the user's commands. Program references appear to play only a minor role. The user command sequence is an important factor in the reference behavior of an interactive session.

Parallel encrypted array multipliers *by S. Vassiliadis, M. Putrino, and E. M. Schwarz, p. 536.* An algorithm for direct two's-complement and sign-magnitude parallel multiplication is described. The partial product matrix representing the multiplication is converted to an equivalent matrix by encryption. Its reduction, producing the final result, needs no specialized adders and can be added with any parallel array addition technique. It contains no negative terms and no extra "correction" rows; in addition, it produces the multiplication with fewer than the minimal number of rows required for a direct multiplication process.

A many-valued logic for approximate reasoning *by S. Di Zenzo, p. 552.* A new system for many-valued logic, the Extended Post system of order p, $p \geq 2$, is proposed as a system of logic supporting reasoning with facts and rules which are reliable to a specified extent. In an Extended Post system there are as many operations of logical disjunction and logical conjunction as there are truth values. The truth value associated with a particular operation of disjunction (conjunction) acts as a threshold value controlling the behavior of the operation. The availability of an extended set of logical operations provides improved flexibility in the symbolic translation of sentences from the ordinary word-language. Extended Post systems are equipped with a semantics in which graded rather than crisp sets correspond to predicates. The system exhibits a "rich" algebraic structure. The p operations of disjunction form a distributivity cycle. To each disjunction there corresponds a dual operation of conjunction, the two operations being distributive to one another. The p conjunctions form a dual distributivity cycle. Both propositional calculus and first-order predicate calculus of *EP* systems are developed. The application to approximate reasoning is described. It is shown that there exist distinct isomorphic copies of fuzzy logic, each corresponding to a distinct level of approximation and being complete to resolution.

Volume 32, Number 5, 1988

Ion transport through protective polymeric coatings exposed to an aqueous phase *by H. Leidheiser, Jr. and R. D. Granata, p. 582.* This paper describes the status of work in our laboratory to develop an improved understanding of the chemical and physical aspects of ion transport through polymeric coatings which are exposed to an aqueous phase.

Characterization of PdSn catalysts for electroless metal deposition *by E. J. M. O'Sullivan, J. Horkans, J. R. White, and J. M. Roldan, p. 591.* A set of electrochemical techniques has been developed to measure the component concentrations and the catalytic activity of the PdSn seeder solutions used to activate insulating substrates for the electroless deposition of Cu. The concentration of Sn(II) was calculated from the limiting current for Sn(II) oxidation, that of Sn(IV) from the difference between the Sn metal-deposition limiting current and the Sn(II) limiting current. The palladium concentration was determined by a stripping analysis after Pd deposition from an oxidized seeder solution. The catalytic activity of the PdSn catalyst was estimated by measuring its activity for the electro-oxidation of formaldehyde (the reducing agent used in the electroless Cu bath) or by the cyclic voltammetric response of a seeded electrode in an inert electrolyte. The cyclic voltammetric technique and transmission electron microscopy examination were used to evaluate various accelerating solutions used to increase the activity of the seeder.

Chemistry at interfaces: Electropositive metals on polymer surfaces *by C. A. Kovac, J. L. Jordan-Sweet, M. J. Goldberg, J. G. Clabes, A. Viehbeck, and R. A. Pollak, p. 603.* This paper presents a study of chemical interactions between polymer surfaces and metal atoms deposited from the vapor phase. Such interactions may play an important role in interfacial metal–polymer adhesion. The chemical nature of the interface formed when an electropositive metal (chromium or cesium) is deposited onto the surface of PMDA–ODA polyimide has been investigated using chemical model studies coupled with photoelectron spectroscopic techniques. X-ray

photoelectron spectroscopy, synchrotron-radiation-excited core-level photoemission, and near-edge X-ray absorption spectroscopy were used to analyze changes in polymer surfaces during deposition of chromium and cesium. Chemical model studies using cyclic voltammetry and UV-visible spectroscopy were performed using several simpler polymers or monomeric model compounds which contained structural subunits of the polyimide. Results of these experiments show that chromium (and other electropositive metals studied so far) initially reacts rapidly with the carbonyl groups of polyimide, causing reduction of the dianhydride portion of the polymer, with concomitant chromium oxidation. Continued deposition of chromium onto the reacted polymer surface results in the formation of chromium carbide, oxide, and nitride species, indicating a disruption of the polymer chemical structure.

Characterization of a bis-maleimide triazine resin for multilayer printed circuit boards *by J. T. Gotro and B. K. Appelt, p. 616.* The thermosetting resin investigated in this study was a mixture of bis-maleimide and bis-cyanate, frequently referred to as BT (bis-maleimide triazine). For printed circuit board applications, a brominated epoxy resin was blended with BT to impart flame resistance. Resin curing was extensively investigated using a combination of thermoanalytical techniques (thermal analysis, heated-cell infrared spectroscopy, dynamic mechanical analysis, and microdielectrometry). Differential scanning calorimetry indicated a minimum of two separate reactions. Fourier-transform infrared spectroscopy provided more detailed information on the cross-linking reactions during the curing. The onset of cyclotrimerization was found to appear at 150°C, correlating with one of the peaks observed in the differential scanning calorimetry measurements. Dynamic mechanical methods were used to investigate the viscosity profile during simulated lamination temperature profiles. Microdielectrometry performed simultaneously with parallel-plate rheometry provided further insight into the physical changes that occur during lamination.

Improvement of adhesion of copper on polyimide by reactive ion-beam etching *by A. L. Ruoff, E. J. Kramer, and C.-Y. Li, p. 626.* In this paper we describe the effect of oxygen-reactive ion-beam etching of a polyimide film to enhance its adhesion to an overlying, subsequently deposited copper film. The adhesion strength of the copper to the polyimide could be increased by as much as a factor of 25 as a result of the etching. Near the etching condition which resulted in optimum strength, the failure mode at the polyimide/copper interface changed from adhesive failure to tensile failure. The latter occurred at the "roots" of a "grass-like" surface structure of the ion-etched polyimide film.

Developer-induced debonding of photoresist from copper *by A. L. Ruoff, E. J. Kramer, and C.-Y. Li, p. 631.* We describe the debonding of a polymeric photoresist film bonded to a thin copper substrate as a result of the diffusion of an organic penetrant into the polymer. The diffusion profile (measured by Rutherford backscattering spectroscopy) consisted of a uniformly swollen layer behind a sharp front which propagated into the polymer at a uniform velocity. Debonding always occurred when the front had penetrated about 12 μm into the polymer (about 1/5 its thickness). The debonding was driven by the release of elastic strain energy created by the swelling.

Physical limits to the useful packaging density of electronic systems *by R. F. Pease and O.-K. Kwon, p. 636.* Increasing the density of electronic circuits and systems has been a major thrust for many years; the benefits are increased speed, reduced power-delay product, and reduced cost. Most of this effort has been directed toward the chip, but during the last decade system performance has been increasingly limited by packaging, and so emphasis has been shifting in that direction. Initially it was believed that heat dissipation was a serious fundamental limit, but advances in heat-sink technology have effectively eliminated that concern. One of the most serious problems is signal distribution. Although we can fabricate submicron metal lines, such lines are not normally practical as chip-to-chip interconnections because their resistance leads to undue signal delay and distortion; increasing their aspect ratio will increase cross talk. It is not clear what constitutes an optimal configuration, but for metals at room temperature a signal-line pitch of 30 to 40 μm appears practical. For low temperatures, and especially for superconducting lines, the pitch could be made very much finer, leading to greatly improved system density.

Electrical design of signal lines for multilayer printed circuit boards *by C. S. Chang, p. 647.* Key aspects of the electrical design of signal lines for multilayer printed circuit boards used in computers are examined. Illustrative calculations are carried out for several signal-line configurations, and associated means are presented for selecting design trade-offs regarding cross talk and skin-effect-induced delay.

Delocalized bonding at the metal–polymer interface *by P. S. Ho, B. D. Silverman, R. A. Haight, R. C. White, P. N. Sanda, and A. R. Rossi, p. 658.* This paper summarizes our current understanding of the nature of the chemical bond formed at the interface between a deposited metal atom and an underlying polyimide surface. The approach in these studies is based on the use of quantum chemical calculations to interpret photoemission spectroscopy results. By focusing on the initial reaction between a chromium atom and the PMDA–ODA polyimide repeat unit, the bonding is demonstrated to be delocalized, arising from the formation of a charge-transfer complex between the metal atom and the PMDA unit of the polyimide. Stabilization of the complex involves the transfer of electronic charge from the metal d states of chromium to the lowest unoccupied molecular orbital of the π system of the PMDA unit of the polyimide. The complex proposed is energetically favored over that involving

a direct local interaction between the chromium atom and one of the carbonyl functional groups. The distribution of single-particle electron energy levels deduced from molecular-orbital calculations can account for the spectroscopy results. The formation of such delocalized metal–polymer complexes is also inferred from a related study of the chromium/PMDA–PDA interface.

Surface analysis and characterization of large printed-circuit-board circuitization process steps *by D. J. Auerbach, C. R. Brundle, and D. C. Miller, p. 669.* We describe our use of surface-analysis techniques to characterize problems encountered in 1980–1981 in the fabrication of large printed circuit boards for the IBM 3081 processor unit. XPS, AES, SAM, SEM, and optical microscopy techniques were used. The two major areas addressed were (a) corrosion at a photoresist/Cu foil interface during electroless Cu plating of circuit lines which resulted in defects in subsequently formed Cu lines, and (b) surface-chemical aspects of a "single-seed" colloidal Pd/Sn catalytic initiation of electroless Cu plating onto epoxy surfaces. The corrosion mechanism responsible for the line defects was identified, and corrective actions suggested. Changes in *surface* composition (Pd/Sn ratio), and *surface* chemical state (Pd^0/Pd^{2+}, Sn^0 $Sn^{2+,4+}$) as a function of process step were correlated with plating effectiveness and led to a means of increasing the surface Pd^0/Sn ratio by as much as an order of magnitude.

Elastic and viscoelastic behavior of a magnetic recording tape *by B. S. Berry and W. C. Pritchet, p. 682.* The mechanical behavior of a trilayer Mylar-based magnetic recording tape has been studied by three complementary methods, applied either to the complete tape or to samples prepared by the selective removal of its front or back coatings. One method provided tensile stress–strain and creep data, another exploited the phenomenon of thermal curling, and a third or mandrel method was used to measure relaxation and recovery in simple bending. Despite the large relative thickness of the Mylar substrate, both the initial stiffness and subsequent relaxation behavior of the tape were strongly influenced by the surface magnetic coatings, and particularly by the oriented and calendered frontcoat, which exhibited elastic anisotropy and an enhanced longitudinal Young's modulus of up to five times that of the Mylar core. As a consequence, the magnetically active frontcoat emerged as the most highly stressed component of the tape, and initially supported almost half of an imposed tensile load. The high initial modulus of the oriented and calendered frontcoat was attributed to the reinforcement provided by the magnetic oxide dispersed in the polymeric frontcoat binder. The substantial viscoelastic behavior of the coatings was also linked to their composite structure, and specifically to the ability of the binder to relax the enhanced initial modulus conferred by the presence of the oxide.

Volume 32, Number 6, 1988

An overview of the basic principles of the Q-Coder adaptive binary arithmetic coder *by W. B. Pennebaker, J. L. Mitchell, G. G. Langdon, Jr., and R. B. Arps, p. 717.* The Q-Coder is a new form of adaptive binary arithmetic coding. The binary arithmetic coding part of the technique is derived from the basic concepts introduced by Rissanen, Pasco, and Langdon, but extends the coding conventions to resolve a conflict between optimal software and hardware implementations. In addition, a robust form of probability estimation is used in which the probability estimate is derived solely from the interval renormalizations that are part of the arithmetic coding process. A brief tutorial of arithmetic coding concepts is presented, followed by a discussion of the compatible optimal hardware and software coding structures and the estimation of symbol probabilities from interval renormalization.

Optimal hardware and software arithmetic coding procedures for the Q-Coder *by J. L. Mitchell and W. B. Pennebaker, p. 727.* The Q-Coder is an important new development in arithmetic coding. It combines a simple but efficient arithmetic approximation for the multiply operation, a new formalism which yields optimally efficient hardware and software implementations, and a new form of probability estimation. This paper describes the concepts which allow different, yet compatible, optimal software and hardware implementations. In prior binary arithmetic coding algorithms, efficient hardware implementations favored ordering the more probable symbol (MPS) above the less probable symbol (LPS) in the current probability interval. Efficient software implementation required the inverse ordering convention. In this paper it is shown that optimal hardware and software encoders and decoders can be achieved with either symbol ordering. Although optimal implementation for a given symbol ordering requires the hardware and software code strings to point to opposite ends of the probability interval, either code string can be converted to match the other exactly. In addition, a code string generated using one symbol-ordering convention can be inverted so that it exactly matches the code string generated with the inverse convention. Even where bit stuffing is used to block carry propagation, the code strings can be kept identical.

Probability estimation for the Q-Coder *by W. B. Pennebaker and J. L. Mitchell, p. 737.* The Q-Coder is an important new development in binary arithmetic coding. It combines a simple but efficient arithmetic approximation for the multiply operation, a new formalism which yields optimally efficient hardware and software implementations, and a new technique for estimating symbol probabilities which matches the performance of any method known. This paper

describes the probability-estimation technique. The probability changes are estimated solely from renormalizations in the coding process and require no additional counters. The estimation process can be implemented as a finite-state machine, and is simple enough to allow precise theoretical modeling of single-context coding. Approximate models have been developed for a more complex multi-rate version of the estimator and for mixed-context coding. Experimental studies verifying the modeling and showing the performance achieved for a variety of image-coding models are presented.

Software implementations of the Q-Coder *by J. L. Mitchell and W. B. Pennebaker, p. 753.* The Q-Coder is an important new development in arithmetic coding. It combines a simple but efficient arithmetic approximation for the multiply operation, a new formalism which yields optimally efficient hardware and software implementations, and a new technique for estimating symbol probabilities which matches the performance of any method known. This paper describes implementations of the Q-Coder following both the hardware and software paths. Detailed flowcharts are given.

A multi-purpose VLSI chip for adaptive data compression of bilevel images *by R. B. Arps, T. K. Truong, D. J. Lu, R. C. Pasco, and T. D. Friedman, p. 775.* A VLSI chip for data compression has been implemented based on a general-purpose adaptive binary arithmetic coding (ABAC) architecture. This architecture permits the reuse of *adapter* and arithmetic *coder* logic in a universal way, which together with application-specific *model* logic can create a variety of powerful compression systems. The specific version of the *adapter/coder* used herein is the "Q-Coder," described in various companion papers. The hardware implementation is in a single HCMOS chip, to maximize speed and minimize cost. The primary purpose of the chip is to provide superior data compression performance for bilevel image data by using conditional binary source models together with adaptive arithmetic coding. The coding scheme implemented is called the Adaptive Bilevel Image Compression (ABIC) algorithm. On business documents, it consistently outperforms such nonadaptive algorithms as the CCITT Group 4 (T.6) Standard and comes into its own when adapting to documents scanned at different resolutions or which include significantly different data such as digital halftones. The multi-purpose nature of the chip allows access to internal partition combinations such as the "Q" adapter/coder, which in combination with external logic can be used to realize hardware for other compression applications. On-chip memory limitations can also be overcome by the addition of external memory in special cases. Other options include the uploading and downloading of adaptive statistics and choices to encode or decode, with or without adaptation of these statistics.

From the fractal dimension of the intermiss gaps to the cache-miss ratio *by D. Thiébaut, p. 796.* This work extends a model proposed by Voldman, Mandelbrot, et al. on the fractal nature of the gaps separating cache misses, and shows how the fractal dimension of the gap distribution can be used to predict the miss ratio experienced by the program that has generated the series of cache misses. This result supports the thesis that the fractal dimension of the distribution of the intermiss gaps is a potentially powerful measure for program characterization.

The kinetics of fast steps on crystal surfaces and its application to the molecular beam epitaxy of silicon *by R. Ghez and S. S. Iyer, p. 804.* Crystal growth by molecular beam epitaxy (MBE) occurs under conditions of high supersaturation. The classic growth theory of Burton, Cabrera, and Frank (BCF) is based on the assumption that surface steps move slowly. Consequently, it requires modifications to be applicable to MBE because the velocities of surface steps may be large. In addition, because such steps are asymmetric structures, as observed experimentally by field ion microscopy, capture probabilities from above and from below a step must differ markedly. Hence the adatom concentration distribution cannot be at equilibrium at steps; there, it also suffers a discontinuity. We propose a model that treats surface step motion as a Stefan problem and that also respects its physical asymmetry. Calculations are presented which extend and complete recently published results that had imposed the restrictive condition of local equilibrium at steps. Step velocity is estimated as a function of supersaturation, degree of asymmetry, and step density. Concentration profiles are then computed; they are found to be generally skewed. In all cases, we show that the behavior of the growing crystal is convective rather than diffusive when the supersaturation is large. Consequently, we can understand the extraordinary insensitivity of the MBE of Si to changes in growth temperature and orientation.

Volume 33, Number 1, 1989

The dynamics of slider bearings during contacts between slider and disk *by R. C. Benson, C. Chiang, and F. E. Talke, p. 2.* The dynamics of a "mini-Winchester" magnetic recording slider are studied during contacts with a hard, rotating memory disk using numerical simulation. An on-line solution of the Reynolds equation is used to calculate the air-film pressure, and a "coefficient-of-restitution" model is used to describe intermittent slider/disk contacts. Studies are made to identify system configurations which reduce the possibility of a "head crash" during contact start/stop.

Functional cache chip for improved system performance *by R. E. Matick, p. 15.* The use of a cache to improve the performance of computing systems is becoming very pervasive, from microprocessors to high-end systems. The general approach has traditionally been to use ordinary fast RAM chips and interface these close to the processor for speed. However, this is far from the ideal solution. The stringent and often conflicting requirements on the cache bandwidth for servicing the processor and minimizing reload time can severely limit attainable performance. The cache need not be the performance-limiting factor if a properly integrated functional cache chip is used. This paper defines the basic requirements of a cache subsystem and shows how these have been or could be implemented in typical systems. Subsequently, the functional requirements of an optimal cache chip design are presented and illustrated.

Modeling electromagnetic interference properties of printed circuit boards *by C. R. Paul, p. 33.* The mathematical modeling of a printed circuit board (PCB) for the prediction of its electromagnetic interference (EMI) properties is investigated. Two key aspects examined are crosstalk and the high-frequency voltage developed between the ends of a PCB land (ground drop). The notion of partial inductance as opposed to loop inductance is the key to predicting the high-frequency voltages that are developed between the two ends of a land. Crosstalk predictions are a by-product of the modeling. Experimental results are shown to illustrate the accuracy of the model. A technique for the accurate measurement of the high-frequency voltage developed between two ends of a PCB land is described and explained in terms of partial inductances.

Analysis of block-paging strategies *by W. H. Tetzlaff, M. G. Kienzle, and J. A. Garay, p. 51.* The performance of interactive paging systems in general and Virtual Machine/System Product (VM/SP) systems with the High Performance Option (HPO) in particular depends upon locality of reference. This storage-management dependency, often considered in the context of individual programs, extends in fact to a significant degree across most virtual-machine transactions. This paper investigates strategies to exploit locality of reference at the system level by analyzing page-reference strings gathered from live systems. Alternative strategies are evaluated using trace-driven simulations.

Statistical properties of selected recording codes *by T. D. Howell, p. 60.* Most recording systems encode their data using binary run-length-limited (RLL) codes. Statistics such as the density of 1s, the probabilities of specific code strings or run lengths, and the power spectrum are useful in analyzing the performance of RLL codes in these applications. These statistics are easy to compute for ideal run-length-limited codes, those whose only constraints are the run-length limits, but ideal RLL codes are not usable in practice because their code rates are irrational. Implemented RLL codes achieve rational rates by not using all code sequences which satisfy the run-length constraints, and their statistics are different from those of the ideal RLL codes. Little attention has been paid to the computation of statistics for these practical codes. In this paper a method is presented for computing statistics of implemented codes. The key step is to develop an exact description of the code sequences which are used. A consequence of the code having rational rate is that all the code-string and run-length probabilities are rational. The method is illustrated by applying it to three codes of practical importance: MFM, (2, 7), and (1, 7).

Volume 33, Number 2, 1989

Geometric tolerancing: I. Virtual boundary requirements *by R. Jayaraman and V. Srinivasan, p. 90.* We examine the representation of geometric tolerances in solid-geometric models from the perspective of two classes of functional requirements. The first class deals with positioning of parts with respect to one another in an assembly, and the second with maintaining material bulk in critical portions of parts. Both are directly relatable to the geometry of the parts. Through examples, we demonstrate that these functional requirements can be captured in a specific form of tolerances designated as virtual boundary requirements (VBRs). We further demonstrate that the only proposed theory of tolerances in solid models, and the current dimensioning and tolerancing standards in industrial practice, are both inadequate for dealing with VBRs. Accordingly, we develop a theoretical basis for the rigorous statement and interpretation of VBRs.

Geometric tolerancing: II. Conditional tolerances *by V. Srinivasan and R. Jayaraman, p. 105.* In a companion paper, we examined the representation of geometric tolerances in solid models from the perspective of certain functional requirements. We showed that assembly and material bulk requirements can be specified as virtual boundary requirements (VBRs). Here, we study the related issue of deriving equivalent alternative specifications. Specifically, we first explore the reasons for converting VBRs to another form of tolerances designated as conditional tolerances (CTs). We then develop a theoretical basis for converting VBRs to CTs and derive CTs for some common and practical VBRs.

We thereby demonstrate the difficulties in finding a general-purpose algorithm for such conversions and also show that some of the CT formulas used in current practice are incorrect.

Conjugate-gradient subroutines for the IBM 3090 Vector Facility *by G. R. di Brozolo and M. Vitaletti, p. 125.* This paper describes a set of optimized subroutines for use in solving sparse, symmetric, positive definite linear systems of equations using iterative algorithms. The set has been included in the Engineering and Scientific Subroutine Library (ESSL) for the IBM 3090 Vector Facility (VF). The subroutines are based on the conjugate-gradient method, preconditioned by the diagonal or by an incomplete factorization. They make use of storage representations of sparse matrices that are optimal for vector implementation. The ESSL vector subroutines are up to six times faster than a scalar implementation of the same algorithm.

Lattice-gas hydrodynamics on the IBM 3090 Vector Facility *by S. Succi, D. d'Humières, and F. Szelényi, p. 136.* After a brief review of the means for characterizing lattice gases using cellular automata rules, we discuss the implementation of the rules for simulating hydrodynamic phenomena which can be described by the Navier–Stokes equations. Special emphasis is placed on data-mapping strategies and implementation through the use of the high speed and large memory resources offered by vector multiprocessors such as the IBM 3090 Vector Facility. We present performance data which pertain to square and hexagonal lattice gases, and discuss the limits of the approach used and its potential extendability to other areas.

A method for generating weighted random test patterns *by J. A. Waicukauski, E. Lindbloom, E. B. Eichelberger, and O. P. Forlenza, p. 149.* A new method for generating weighted random patterns for testing LSSD logic chips and modules is described. Advantages in using weighted random versus either deterministic or random test patterns are discussed. An algorithm for calculating an initial set of input-weighting factors and a procedure for obtaining complete stuck-fault coverage are presented.

Large-area fault clusters and fault tolerance in VLSI circuits: A review *by C. H. Stapper, p. 162.* Fault-tolerance techniques and redundant circuits have been used extensively to increase the manufacturing yield and productivity of integrated-circuit chips. Presented here is a review of relevant statistical models which have been used to account for the effects on manufacturing yield of the large-area defect and fault clusters commonly encountered during chip fabrication. A statistical criterion is described for determining whether such large-area clusters are present.

Small-area fault clusters and fault tolerance in VLSI circuits *by C. H. Stapper, p. 174.* In previous treatments of the manufacturing yield of fault-tolerant integrated-circuit chips, fault clusters were either assumed to be absent or relatively large in area. Presented here is a treatment in which the occurrence of small-area fault clusters is assumed. Four different types of statistical distributions are considered, and a criterion is described for determining whether small-area fault clusters are present.

Translating object specifications into a computer-generated three-dimensional graphic to be reproduced as a high efficiency, reflection photo-polymer hologram suitable for mass-production *by W. J. Molteni, Jr. and D. Small, p. 178.* A process is described for translating the specifications of an object and its interrelationship with another object into a three-dimensional computer graphic and then into a photo-polymer hologram. The capability to translate specifications about objects and their interrelationships into accurate holograms without having to create either a physical model or the manufactured object itself opens exciting possibilities in the areas of creative design and communication. It may assist in the manufacturing process by allowing designers to specify objects and to study accurate, three-dimensional representations of those specifications, including interrelationships with other objects, without the need for an actual physical model. Finally, the hologram may become a means for effective representation of an image produced through the use of three-dimensional computer graphics for people without access to appropriate computer graphics. The process described here was divided into two segments. The MIT Media Lab was responsible for creating a sequence of computer-generated images and transferring those images to film. Polaroid Corporation was responsible for creating the hologram from the images on film.

Volume 33, Number 3, 1989

Preface *by K. A. Müller and J. G. Bednorz, p. 199.*

Current understanding of electronic structure and some difficulties with cuprate semiconductors *by K. Kitazawa, p. 201.* Experimental observations of the high-T_c cuprate superconductors are reviewed from the perspective of electronic structures. On the basis of Mott–Hubbard-type band splitting, the semiconductivity, antiferromagnetism, and metallic nature of the cuprate oxides are discussed as a function of the dopant concentration. Then the involvement of the O $2p$ band which falls between the lower and upper Hubbard bands, as determined by electron spectroscopy, is discussed. Finally, the complex nature of the Fermi surface is considered, and the importance of the involvement of both the O $2p$ and the lower Hubbard band is stressed.

High-T_c superconductivity in bismuthates—How many roads lead to high T_c? *by B. Batlogg, R. J. Cava, L. F. Schneemeyer, and G. P. Espinosa, p. 208.* The superconducting transition temperature in $BaBiO_3$-based superconductors exceeds 30K. Magnetic measurements are analyzed to give their density of states at E_F, $N^*(0) \propto \gamma$. The

uniqueness of Bi–O and Cu–O superconductors is revealed in an updated T_c–γ plot. The two classes of compounds share basic electronic properties, particularly a partially unoccupied band with significant O 2*p* character, which might favor a common pairing mechanism.

Interim report on the charge-transfer resonance model for the Cu–O superconductors *by C. M. Varma, p. 215.* The two-band model including intra-atomic repulsion on Cu and near-neighbor Cu–O repulsion for the high-temperature superconductors is supported by a variety of experiments and theoretical calculations as the minimum necessary. The specific mechanism for high T_c through charge-transfer resonances was proposed because the known alternative mechanisms—phonons and magnetic excitations—were believed unlikely. The case for charge-transfer-resonance-induced high T_c has, however, not yet been proven. Equally important, the various anomalies in the metallic state are not yet understood. However, calculations on the model do show a charge-transfer-gap insulating state which is antiferromagnetic at and near 1/2 filling, a metallic state for intermediate filling with effective particle–particle attraction, and a charge-transfer instability beyond a certain filling.

Oxygen "disorder" and the structures of high-T_c superconductors by neutron powder diffraction *by A. W. Hewat, E. A. Hewat, P. Bordet, J.-J. Capponi, C. Chaillout, J. Chenavas, J.-L. Hodeau, M. Marezio, P. Strobel, M. François, K. Yvon, P. Fischer, and J.-L. Tholence, p. 220.* All of the high-T_c perovskite superconductors appear to show disorder of certain oxygen atoms. In $(La,Sr)_2CuO_4$ and perhaps also in $YBa_2Cu_3O_7$ this is associated with a structural transition. The Bi and Tl superconductors, for which we now have neutron structural data on four different phases, also show oxygen "disorder" which may be associated with valence fluctuations. In $Tl_2Ba_2CuO_6$, electron holes are created by the absence of 1/8 of the atoms in the TlO plane, producing a marked superstructure. However, this material is not superconducting if the superstructure is well ordered, with an orthorhombic (strictly monoclinic) structure. The T_c appears to depend on the disorder of the superstructure to produce a pseudotetragonal metric in which the oxygen atoms within the TlO plane are distributed over four equivalent sites about the center of the Tl square.

Tl–Ca–Ba–Cu–O superconducting oxides *by R. B. Beyers, S. S. P. Parkin, V. Y. Lee, A. I. Nazzal, R. J. Savoy, G. L. Gorman, T. C. Huang, and S. J. La Placa, p. 228.* This paper reviews structural studies of Tl–Ca–Ba–Cu–O superconducting oxides by the authors and others and points out directions for future work.

Infrared studies of the normal and superconducting states of $Y_1Ba_2Cu_3O_7$ *by R. T. Collins, Z. Schlesinger, F. H. Holtzberg, P. Chaudhari, and C. A. Feild, p. 238.* We describe infrared measurements of the *a*–*b*-plane response of $Y_1Ba_2Cu_3O_7$ crystals with $T_c \approx 92$K. We observe a self-energy structure at a characteristic energy of 500 cm^{-1} ($8kT_c$), the appearance of which coincides with the transition to the superconducting state. The nature of this self-energy anomaly is consistent with its identification as a nodeless *a*–*b*-plane energy gap at $2\Delta \approx 8kT_c$. On the basis of temperature-dependent measurements above T_c, we suggest that the normal state can be primarily characterized by a carrier band with an enhanced low-frequency mass and a frequency-dependent scattering rate. Our data indicate that these arise from coupling to an excitation spectrum with characteristic frequencies up to $\omega_c \sim 700$ cm^{-1} and a coupling strength of $\lambda \approx 2$–3.

On the nature of high-temperature superconductivity *by V. J. Emery, p. 246.* A picture of the electronic structure, magnetism, and superconductivity in high-T_c oxides is obtained from a simple analysis of experiments and models of the copper oxide planes. It is shown that magnetism is associated with holes on copper and superconductivity with holes on oxygen. The pairing force is not retarded. Questions about the motion of charges in an antiferromagnetic background and the many-body theory of high-temperature superconductivity are discussed. Differences between the cuprates and doped $BaBiO_3$ are emphasized.

Measurements on thin-film high-T_c superconductors *by A. Kapitulnik and K. Char, p. 252.* We report on the fabrication and properties of high-T_c superconducting films, and discuss the possible origin of the linear resistivity. We further discuss the *c*-axis conductivity. We present data to show that the coherence length is 12 Å parallel to the plane and 2 Å perpendicular to the plane in the superconducting state. Tunneling data show that the energy gap is large, with $2\Delta/kT_c \approx 7$.

Granular Josephson and quantum interference effects in HTC ceramic superconductors *by C. E. Gough, p. 262.* Josephson effects in high-T_c superconductors are briefly reviewed, with specific reference to granular ceramic materials and SQUID device applications. It is suggested that the inductance associated with intergranular current loops may play an important role, even in determining the bulk superconductivity properties, as in weak-link superconducting rings. Evidence for quantum interference effects within intergranular current loops is presented. In ultra-low fields, the observed temperature dependence of thermally activated flux creep cannot be described by a simple granular superconductor model of equally spaced pinning centers, but would seem to imply a hierarchy of pinning sites of variable strength. The development of liquid-nitrogen-cooled rf and dc SQUIDs is described, and the noise levels currently achieved are presented.

Quasi-elastic and inelastic neutron-scattering studies of superconducting $La_{2-x}Sr_xCuO_4$ *by R. J. Birgeneau, Y. Endoh, Y. Hidaka, K. Kakurai, M. A. Kastner, T. Murakami, G.*

Shirane, T. R. Thurston, and K. Yamada, p. 270. We review the results of recent neutron-scattering studies of the spin fluctuations in samples of $La_{1.89}Sr_{0.11}CuO_4$ which are ~80% superconducting with $T_c = 10K$. The structure factor, $S(\vec{Q})$, reflects three-dimensional modulated spin correlations with an in-plane correlation length of order 18 ± 6 Å. The fluctuations evolve with temperature from being predominantly dynamic at high temperatures to mainly quasi-elastic ($|\Delta E| < 0.5$ meV) at low temperatures. No significant differences are observed in the normal and superconducting states.

NMR study of magnetism and superconductivity in high-T_c oxides *by Y. Kitaoka, K. Ishida, K. Fujiwara, K. Asayama, H. Katayama-Yoshida, Y. Okabe, and T. Takahashi, p. 277.* The results of ^{139}La and ^{63}Cu nuclear quadrupole resonance (NQR) on La–Ba,Sr–Cu–O, Y–Ba–Cu–O, and Bi–Pb–Sr–Ca–Cu–O compounds and of ^{17}O nuclear magnetic resonance (NMR) on Y–Ba–Cu–O are extensively reviewed. As for the magnetism, the phase diagram for the La system studied by a ^{139}La NQR experiment is presented, with evidence of the disordered magnetic state between the 3D-antiferromagnetic (AF) ordered state and the superconducting state. With respect to its superconducting nature, the nuclear spin-lattice relaxation behavior (T_1) of Cu in the CuO_2 plane has been found to be unconventional above and below T_c for all compounds, with no signatures expected for a nonmagnetic metal and a BCS superconductor, respectively. The behavior of T_1 of Cu above T_c is shown to be dominated by AF fluctuation of Cu *d* spins. In contrast, an enhancement of $1/T_1$ of ^{17}O has been observed just below T_c, which is similar to a BCS case.

Magnetic frustration model and superconductivity in doped planar CuO_2 systems *by A. Aharony, R. J. Birgeneau, and M. A. Kastner, p. 287.* We present a model for the magnetic phases and superconductivity in doped planar CuO_2 systems. Electronic holes on the oxygen ions introduce local ferromagnetic exchange couplings between the Cu spins. The resulting frustration destroys the antiferromagnetic state characterizing the undoped planes, and generates a new spin-glass phase. This frustration also yields an attractive interaction between the holes, whose range decreases with increasing doping. We use the BCS approximation to obtain an excellent estimate of the superconducting transition temperature $T_c(x)$ for $La_{2-x}Sr_xCuO_4$.

Short-coherence-length superconductors *by G. Deutscher, p. 293.* The new high-T_c oxides present some anomalous electromagnetic properties, such as low critical current densities, a reversible behavior of the magnetization at fields much lower than H_{c2}, and internal Josephson effects, that distinguish them from the conventional low-T_c metals and alloys. These anomalous properties were first observed in bulk-sintered samples and were often ascribed to the poor connectivity of these ceramics. More recently, a qualitatively similar behavior has been observed in single crystals and oriented films. The fundamental role of the short coherence length in determining the behavior of the high-T_c oxides is discussed. We show that the short coherence lengths at the local depressions of the order parameter at crystallographic defects lead to reduced critical currents and cause glassy behavior in the vicinity of T_c.

Critical current measurements in single crystals and single-grain boundaries in $YBa_2Cu_3O_7$ films *by P. Chaudhari, D. Dimos, and J. Mannhart, p. 299.* The temperature, magnetic field, and orientation dependence of the critical current density of superconducting $YBa_2Cu_3O_7$ films have been determined from transport measurements. The results support a model of flux creep within single grains and weak-link coupling across grain boundaries.

Glassy behavior of high-T_c superconductors *by I. Morgenstern, p. 307.* This paper deals with the question of flux creep or glassy behavior in high-T_c superconducting single crystals. It is shown that the flux creep picture is merely a phenomenological approach to the glassy behavior for relatively short times and low temperatures. Glassy effects are predicted for temperatures between 70% and 95% T_c and magnetic fields in the range of 0.03 T to 0.2 T. The glass concept can be understood as a generalization of the traditional flux creep picture. A hierarchy of energy barriers dominates the physical behavior. An important technical aspect is the influence of the glassiness on critical currents.

Muon-spin rotation experiments in high-T_c superconductors and related materials *by H. Keller, p. 314.* Recent muon-spin rotation μSR experiments in high-T_c superconductors and related antiferromagnetic materials are reviewed. The possibilities and the limitations of the μSR method for investigating these materials are briefly discussed. In a high-T_c superconductor, μSR is an ideal tool with which to study the local magnetic field distribution at the muon site, allowing a determination of the London penetration depth. It is further shown that μSR experiments may contribute to the microscopic understanding of the superconducting glass state in the high-T_c oxides. In the related antiferromagnetic materials μSR is a sensitive method for detecting frozen local magnetic moments.

Low-field microwave absorption in single-crystal superconducting $YBa_2Cu_3O_{7-\delta}$ *by K. W. Blazey and F. H. Holtzberg, p. 324.* The low-field microwave absorption line spectrum of single crystal of superconducting $YBa_2Cu_3O_{7-\delta}$ has been studied as a function of the external magnetic field. The threshold microwave power necessary to nucleate fluxons is found to vary with field in such a way that only about one thousandth of the junction length is active in interacting with the microwaves to create fluxons.

Memory effects in $YBa_2Cu_3O_{7-\delta}$ single crystal *by C. Rossel, Y. Maeno, and F. H. Holtzberg, p. 328.* Measurements of the time dependence of zero-field-cooled (ZFC) and field-cooled

(FC) magnetization M in $YBa_2Cu_3O_{7-\delta}$ single crystal have been performed as a function of temperature and magnetic field. The appearance of an echo-like feature in the decay rate $S = dM/d \ln t$ at an observing time t equal to the waiting time t_w during which the specimen was prepared at a given field H_0 and temperature T reveals aging effects in the superconducting state. Similar phenomena reported in spin glasses seem to validate the picture of a superconducting glass state in $YBa_2Cu_3O_{7-\delta}$.

Positron annihilation and high-temperature superconductivity *by M. Peter, p. 333.* After a brief review of the theory of positron annihilation techniques and of experimental principles, we give examples of the successful determination of electron momentum density and Fermi surfaces in alkali metals, transition elements and compounds, and cerium. We then discuss the application of positrons in superconducting oxides. So far the best results have been obtained in $YBa_2Cu_3O_{7-x}$, with confirmation of calculated band structure and observation of discontinuity at the Fermi energy.

Nuclear magnetic resonance in high-T_c superconductors *by M. Mehring, p. 342.* This paper is aimed at the nonspecialist in nuclear magnetic resonance who wants to know what NMR can do to increase his understanding of high-T_c superconductors. Most NMR results are discussed in an illustrative manner to facilitate intuitive understanding. Several specific NMR experiments are presented which demonstrate the variety of this experimental technique. Special emphasis is given to the following aspects: ionic charges and quadrupole interaction, local fields and magnetic ordering, conduction electrons and Knight shifts, quasiparticle excitations, and nuclear spin-lattice relaxation.

Critical temperature and the Ginzburg–Landau theory of layered high-temperature superconductors *by T. Schneider, p. 351.* Using the mean-field approximation, we study a model for quasi-two-dimensional superconductors. The interlayer coupling, assumed to be mediated by a small electron-hopping term, is found to leave T_c practically unaffected. Thus, a three-dimensional pairing mechanism is required to explain the observed rise in T_c with decreasing average layer spacing in the Bi and Tl compounds. Taking the inhomogeneities of intrinsic or extrinsic nature into account, we find, in the dirty limit, corrections to the conventional anisotropic Ginzburg–Landau behavior—an upward curvature of the upper critical fields which appears to be a universal feature of layered superconductors.

Instability and high-T_c superconductivity *by M. Kataoka, p. 356.* The interrelation between instability and a high superconducting transition temperature T_c is argued theoretically in order to construct a model for the origin of the high T_c in the perovskite-type oxides. It is shown that when two nearly degenerate bands overlapping on the Fermi energy ε_F become unstable against spontaneous splitting to give rise to a charge redistribution, effective interactions between two electrons in the same bands become attractive, and the attractive interactions are strongly enhanced by increasing the degree of instability. One cause of this instability is the electron–phonon coupling, which results in lattice-instability-enhanced superconductivity; another cause is the Coulomb interaction between electrons, which results in electron-instability-caused superconductivity. The latter mechanism is successfully applied to the perovskite-type high-T_c superconductors. Some guidelines for obtaining a high T_c are presented on the basis of the present idea.

A review of elastic properties of high-T_c superconductors and some related C_p results *by K. Fossheim and T. Lægreid, p. 365.* First a brief survey is given of what can be learned about important superconducting and normal-state properties by ultrasonic and other elastic measurements. Some of the characteristic elastic properties of the $La_{2-x}(Ba, Sr)_xCuO_4$ and $YBa_2Cu_3O_7$ systems are reviewed. In the La-based family it is shown how the elastic observations are closely related to structural and soft-mode properties. The physics of $YBa_2Cu_3O_7$ is shown to be more complex. Finally, our recent results on the fluctuation contribution to the specific heat near T_c in $YBa_2Cu_3O_7$ are discussed.

Electronic structure studies of high-T_c superconductors by high-energy spectroscopies *by J. Fink, N. Nücker, H. A. Romberg, and J. C. Fuggle, p. 372.* A review of our high-energy spectroscopy studies of the electronic structure of the new high-T_c superconductors is given. X-ray-induced photoelectron spectroscopy, bremsstrahlung–isochromat spectroscopy, Auger electron spectroscopy, and electron energy-loss spectroscopy have been used. Parameters determining the correlated electronic structure have been derived, together with information on the nature and the symmetry of the charge carriers.

A new structural modification of superconducting $La_{2-x}M_xCuO_4$ *by J. D. Axe, D. E. Cox, K. Mohanty, H. Moudden, A. R. Moodenbaugh, Y. Xu, and T. R. Thurston, p. 382.* $La_{1.9}Ba_{0.1}CuO_4$ has been observed to undergo the following sequence of transformations upon cooling from 300K: tetragonal ($I4/mmm$) $\rightarrow$ ($T_0 = 270$K) orthorhombic (Cmca) $\rightarrow$ ($T_1 = 52$K) tetragonal ($P4_2/ncm$). The newly discovered low-temperature tetragonal structure can be understood geometrically as arising from a coherent superposition of the two domain modifications of the orthorhombic structure. Dynamically, it results from a second instability in the twofold degenerate soft modes of the high-temperature tetragonal phase. Energetically, the system can be modeled as an XY-spin system with temperature-dependent quartic anisotropy $v(T)$, and the low-temperature transformation coincides with an isotropic point $v(T_1) = 0$. The relationship of the newly discovered transformation to other anomalous properties and to superconductivity is discussed briefly.

Energy dependence of the Andreev reflection of $YBa_2Cu_3O_{7-\delta}$ *by H. van Kempen, H. F. C. Hoevers, P. J. M. van Bentum, A. J. G. Schellingerhout, and D. van der Marel, p. 389.* Measurements of the energy dependence of the Andreev reflection have been performed on a $Ag–YBa_2Cu_3O_{7-\delta}$ interface. The observation of the Andreev reflection indicates a ground state of zero-momentum pairs. It is shown that, in principle, the bulk Δ (electron pair potential) can be determined from the energy dependence of the Andreev reflection. In the present experiment, however, due to the limited mean free path of the electrons in the silver, only a lower limit of Δ was found.

Volume 33, Number 4, 1989

Modeling and image processing for visualization of volcanic mapping *by M. T. Pareschi and R. Bernstein, p. 406.* In countries such as Italy, Japan, and Mexico, where active volcanoes are located in highly populated areas, the problem of risk reduction is very important. Actual knowledge about volcanic behavior does not allow deterministic event prediction or the forecasting of eruptions. However, areas exposed to eruptions can be analyzed if eruption characteristics can be inferred or assumed. Models to simulate volcanic eruptions and identify hazardous areas have been developed by collaboration between the IBM Italy Pisa Scientific Center and the Earth Science Department of Pisa University (supported by the Italian National Group of Volcanology of the Italian National Research Council). The input to the models is the set of assumed eruption characteristics: the typology of the phenomenon (ash fall, pyroclastic flow, etc.), vent position, total eruptible mass, wind profile, etc. The output of the models shows volcanic product distribution at ground level. These models are reviewed and their use in hazard estimation (compared with the more traditional techniques currently in use) is outlined. Effective use of these models, by public administrators and planners in preparing plans for the evacuation of hazardous zones, requires the clear and effective display of model results. Techniques to display and visualize such data have been developed by the authors. In particular, a computer program has been implemented on the IBM 7350 Image Processing System to display model outputs, representing both volume (in two dimensions) and distribution of ejected material, and to superimpose the displays upon satellite images that show 3D oblique views of terrain. This form of presentation, realized for various sets of initial conditions and eruption times, represents a very effective visual tool for volcanic hazard zoning and evacuation planning.

Natural quadrics: Projections and intersections *by M. A. O'Connor, p. 417.* Geometrical modelers usually strive to support at least solids bounded by the results of Boolean operations on planes, spheres, cylinders, and cones, that is, the natural quadrics. Most often this set is treated as a subset of the set of quadric surfaces. Although the intersection of two quadrics is a mathematically tractable problem, in implementation it leads to complexity and stability problems. Even in the restriction to the natural quadrics these problems can persist. This paper presents a method which, by using the projections of natural quadrics onto planes and spheres, reduces the intersection of two natural quadrics to the calculation of the intersections of lines and circles on planes and spheres. In order to make the claims of the method easily verifiable and provide the tools necessary for implementation, explicit descriptions of the projections are also included.

Femtosecond laser studies of the relaxation dynamics of semiconductors and large molecules *by C. L. Tang, F. W. Wise, M. J. Rosker, and I. A. Walmsley, p. 447.* The use of femtosecond lasers and the related optical correlation spectroscopic technique for studying the relaxation dynamics of semiconductors and photoexcited molecules are reviewed. In particular, the results on the intraband relaxation of nonequilibrium carriers in GaAs and related compounds and quantum well structures are summarized. The optical correlation technique also led to the observation of quantum beats in the femtosecond time domain corresponding to the direct observation of molecular vibrations in the time domain.

Effects of quasiperiodic (Penrose tile) symmetry on the eigenvalues and eigenfunctions of the wave equation *by S. He and J. D. Maynard, p. 456.* In addition to the basic crystalline and amorphous structures for solids, it is possible that solids may also form with a quasiperiodic, or Penrose tile, structure. A current problem in condensed-matter physics is to determine how this structure affects the various physical properties of a material. A fundamental question involves the consequences of quasiperiodic symmetry in the eigenvalue spectrum and eigenfunctions of a wave equation. While rigorous theorems have been derived for one-dimensional systems, there is currently no known "quasi-Bloch theorem" for two and three dimensions. To gain insight into this problem, an acoustic experiment has been used to study a two-dimensional wave system with a Penrose tile symmetry. The results show an eigenvalue spectrum containing bands and gaps with widths which are in the ratio of the Golden Mean, $(\sqrt{5}+1)/2$.

Time series in *M* dimensions: The power spectrum *by R. H. Yetzer, p. 464.* The approach presented here extends the modeling of *M*-dimensional (spatial) time series from the time domain into the frequency domain. The autocovariance function for an *M*-dimensional time series is transformed to obtain the power spectrum in *M* dimensions. The latter describes the variance within the series and can be used to identify dependencies and/or test the adequacy of a fitted model. An example is provided.

Two-level coding for error control in magnetic disk storage products *by A. M. Patel, p. 470.* Error-control coding has played a significant role in the design and development of magnetic recording storage products. The trend toward higher densities and data rates presents continuing demands

for an ability to operate at a lower signal-to-noise ratio and to tolerate an increased number of correctable errors. Heretofore, the magnetic disk storage products used coding schemes that provided correction of one burst of errors in a record of length ranging from a few bytes of data to a full track on the disk. In this paper, we present a new coding architecture that facilitates correction of multiple-burst errors in each record in a typical disk storage application. This architecture embodies a two-level coding scheme which offers high coding efficiency along with a fast decoding strategy that closely matches the requirements of on-line correction of multiple bursts of errors. The first level has a smaller block delay and provides very fast correction of most of the errors commonly encountered in an average disk file. The second level, on a larger block size, provides reserve capability for correcting additional errors which may be encountered in a device with symptoms of a weaker component or an oncoming failure. The new IBM 3380J and 3380K disk files use a two-level scheme that is designed around the coding structure of the extended Reed–Solomon code. This design and the related encoding and decoding methods and implementation are presented in detail.

Volume 33, Number 5, 1989

Frames, semantic networks, and object-oriented programming in APL2 *by M. Alfonseca, p. 502.* This paper discusses the capabilities of APL2 for the implementation of frame systems and semantic networks, and for the use of object-oriented programming techniques. The fact that the frame is a basic data structure of APL2 makes this language very appropriate for the development of artificial intelligence applications using the indicated techniques. Examples are given of the way in which they may be applied to realistic situations.

Dependability evaluation of a class of multi-loop topologies for local area networks *by W. E. Smith and K. S. Trivedi, p. 511.* Local area networks have been developed using both ring and bus topologies. Multi-loop and multi-connected topologies have been proposed to improve the throughput and dependability of single-loop networks. We evaluate the dependability of a class of multi-connected loop topologies called forward loop, backward hop (FLBH) networks and compare them to simple ring networks.

Architecture, design, and operating characteristics of a 12-ns CMOS functional cache chip *by R. Matick, R. Mao, and S. Ray, p. 524.* The architecture, design, and implementation of a high-performance cache require a detailed consideration of the overall system functions closely coupled with the proper mapping and integration of these functions into the circuits and arrays. This approach has resulted in a new cache chip which incorporates a number of unique on-chip functions as well as unique design, providing a one-cycle cache in which translation can be overlapped with cache access. In order to achieve high average performance, a cache should give the appearance of being a two-ported array in order to provide high bandwidth both to the processor during normal execution and to the main memory during reload. But a true two-port design is undesirable as well as unnecessary, especially since the reload process is typically limited by the memory speed and other system parameters. A significant improvement can be obtained by judicious choice as well as proper integration of some critical functions placed directly on the cache chips. This paper describes these functions and integration onto a 72K-bit static RAM chip, implemented in 1-μm CMOS technology for high speed and overall system performance. In addition, the chip I/O is selectable for either ECL or TTL compatibility.

Order in the domain structure in soft-magnetic thin-film elements: A review *by H. A. M. van den Berg, p. 540.* The domain structure and its development in thin plane-parallel soft-magnetic elements have been investigated from both the experimental and the theoretical point of view. The experimental observations for verifying the predictions have been realized by means of the Bitter, Kerr, and Lorentz techniques.

In the first part, a self-consistent domain theory, based on micromagnetic principles, is unfolded for two-dimensional solenoidal magnetization distributions present in ideally soft-magnetic thin-film objects that are rectangular cylinders. The solenoidality implies that both the external field and the conduction currents are taken as zero. Two types of domain structures are distinguished: the basic structures in simply connected regions and the parallel configurations in special types of multiply connected regions—the parallel regions. A decomposition of the area of the object into disjunct subregions, either simply connected or of the parallel type, whose union completely covers the object, is put forward. A procedure for constructing all feasible parallel regions is presented. In each region, the appropriate solenoidal magnetization distribution is specified with which the magnetization M is taken parallel to the boundaries of the subregion. Thus, all the domain structures possible in the thin-film objects with arbitrary lateral shapes can be constructed. A number of experimental examples are provided.

In the second part, the M distribution is studied on a local scale, at which the requirement of solenoidality is dropped; i.e., external fields and conduction currents are allowed. The concept of the domain-wall cluster is introduced in order to obtain the maximum information about the M configuration in the entire object. Here, we employ the fact that domain walls are the preeminently visible features and that most information is available at those locations where a number of these walls meet. A domain-wall cluster is the collection of all domain walls that have one region—the so-called cluster knot—in common. Three different categories of clusters characterized by the positions of their cluster knots with respect to the edges of the thin-film object are distinguished.

Wall clusters with cluster knots at two, one, and no edges are defined as the corner, edge, and free clusters, respectively. General features of the magnetization distribution near the cluster knots are discussed for each of the above classes. The reversible transformations of the clusters are reviewed. Two different types of these conversions are recognized, to wit the cluster creation (fading) and the cluster furcation (fusion). Experimental evidence of these relationships is provided.

In the third part, the domain structures are considered as a concatenation of domain-wall clusters. During the domain-structure transformations, clusters are added to and removed from the domain-wall network. The conversions are reversible along specific branches of the hysteresis curves at which the changes can be comprehended in terms of the above reversible cluster conversions. Notwithstanding the reversible character at these branches, the domain configuration often develops itself into a subminimum of the energy, from which sudden irreversible transformations take place toward other branches with lower energy. In many cases, the latter alterations are attended by jumpwise adaptations in the overall object magnetization component along the field, and reveal themselves in the hysteresis curve. The part of the internal domain-wall structure in the hysteresis is elucidated, and its dependence on the film thickness is emphasized. Many examples are given for the purpose of demonstrating the strongly interwoven character of the domain network, the prehistory in the magnetic sense, the internal structure of the domain walls, and the macroscopic object hysteresis.

Volume 33, Number 6, 1989

Coding for constrained channels: A comparison of two approaches *by P. A. Franaszek, p. 602.* We discuss the relation between some early techniques for constrained channel coding and more recent ones adapted from the mathematical area of symbolic dynamics. A primary difference between the two is that the latter focus on issues of code existence, whereas the former were primarily concerned with code construction and optimality.

Matrix print actuator for dot band printer *by J. L. Zable and E. F. Helinski, p. 609.* The design selection of an impact matrix print actuator is described, along with its method of design and optimization. The dot band printer concept is discussed in conjunction with operational requirements for the hammer; these requirements, and methods of meeting them, are presented.

Preserving the integrity of cyclic-redundancy checks when protected text is intentionally altered *by D. R. Irvin, p. 618.* As a digitally encoded message traverses a series of point-to-point communication links, it may be necessary to change the contents of that message at an intermediate station. If bit errors are introduced by the intermediary while the text is unprotected, these errors will be subsequently undetectable by cyclic redundancy checks. An algorithm is presented here for ensuring that such errors will not go undetected. Since the cyclic redundancy check is based on a linear mathematical operation, the frame-check sequence may be modified, rather than recalculated, by each intermediary changing the protected text. A frame-check sequence constructed in this way will reveal any errors introduced in the transmission path when the message is finally examined at the ultimate destination. Examples of the proposed technique applied to various local-area network bridges are developed. The technique is shown to be beneficial in these examples when the internal bit-error ratio of the text-changing device exceeds 10^{-19} on unprotected paths.

Representing knowledge with functions and Boolean arrays *by K. Fordyce, J. Jantzen, G. A. Sullivan, Sr., and G. A. Sullivan, Jr., p. 627.* Over the past eighteen years a variety of advanced decision support systems have been built with knowledge-based expert system (KBES) components. For the past eight years, a knowledge representation and manipulation (KRM) scheme called FABA (Functions And Boolean Arrays) has been used. It has two basic principles. First, knowledge is viewed as a functional mapping between input and output variables, where the functions are expressed as *fact tables* or *bases* and *procedure modules.* Second, the function network can be represented with Boolean arrays. The basics of FABA, its implementation in APL2, and a simple example of FABA's application in a manufacturing dispatch application for IBM's semiconductor facility in Burlington, Vermont, are described in this paper.

Fault-simulation programs for integrated-circuit yield estimations *by C. H. Stapper, p. 647.* Three programs are described here which have been used for integrated-circuit yield modeling at the IBM facility in Essex Junction, Vermont. The first program generates negative binomial distributions which are used to represent the frequency distribution of the number of faults per chip. Calculations with the generalized combination function *A ! B* in *APL* are limited to simulations of up to 99,999 faults, and can take too much computer time to run. These limitations are eliminated when the calculations make use of the scan function. The second program simulates clustered fault locations on a map. The clusters are initially generated using a radial Gaussian probability distribution. Each fault location is stored as a complex number, which facilitates the use of cluster-shaping programs that are also described. In a third program, another simulator of fault maps, faults are added as a function of time. This program also results in fault distributions that are clustered. In addition, it produces frequency distributions that very closely approximate negative binomial distributions.

Volume 34, Number 1, 1990

Preface *by N. M. Donofrio, p. 2.*

The evolution of RISC technology at IBM *by J. Cocke and V. Markstein, p. 4.* This paper traces the evolution of IBM RISC architecture from its origins in the 1970s at the IBM Thomas J. Watson Research Center to the present-day IBM RISC System/6000® computer. The acronym RISC, for Reduced Instruction-Set Computer, is used in this paper to describe the 801 and subsequent architectures. However, RISC in this context does not strictly imply a reduced number of instructions, but rather a set of primitives carefully chosen to exploit the fastest component of the storage hierarchy and provide instructions that can be generated easily by compilers. We describe how these goals were embodied in the 801 architecture and how they have since evolved on the basis of experience and new technologies. The effect of this evolution is illustrated with the results of several benchmark tests of CPU performance.

The IBM RISC System/6000 processor: Hardware overview *by H. B. Bakoglu, G. F. Grohoski, and R. K. Montoye, p. 12.* A highly concurrent superscalar second-generation family of RISC workstations and servers is described. The RISC System/6000® family is based on the new IBM POWER (Performance Optimization With Enhanced RISC) architecture; the hardware implementation takes advantage of this powerful RISC architecture and employs sophisticated design techniques to achieve a short cycle time and a low cycles-per-instruction (CPI) ratio. The RS/6000 CPU features multiple-instruction dispatch, multiple functional units that operate concurrently, separate instruction and data caches, and zero-cycle branches. In this superscalar implementation, at a given cycle the equivalent of five operations can be executed simultaneously (a branch, a condition-register operation, and a floating-point multiple-add). The RS/6000 family supports the IBM Micro Channel architecture as well as high-speed serial optical links to provide a high-bandwidth I/O subsystem.

IBM RISC System/6000 processor architecture *by R. R. Oehler and R. D. Groves, p. 23.* This paper describes the hardware architecture of the IBM RISC System/6000® processor, which combines basic RISC principles with a partitioning of registers by function into multiple ALUs. This allows a high degree of parallelism in execution and permits a compiler to generate highly optimized code to manage the interaction among parallel functions. Floating-point arithmetic is integrated into the architecture, and floating-point performance is comparable to that of many vector processors.

Machine organization of the IBM RISC System/6000 processor *by G. F. Grohoski, p. 37.* The IBM RISC System/6000® processor is a second-generation RISC processor which reduces the execution pipeline penalties caused by branch instructions and also provides high floating-point performance. It employs multiple functional units which operate concurrently to maximize the instruction execution rate. By employing these advanced machine-organization techniques, it can execute up to four instructions simultaneously. Approximately 11 MFLOPS are achieved on the LINPACK benchmarks.

Design of the IBM RISC System/6000 floating-point execution unit *by R. K. Montoye, E. Hokenek, and S. L. Runyon, p. 59.* The IBM RISC System/6000® (RS/6000) floating-point unit (FPU) exemplifies a second-generation RISC CPU architecture and an implementation which greatly increases floating-point performance and accuracy. The key feature of the FPU is a unified floating-point multiply-add-fused unit (MAF) which performs the accumulate operation $(A \times B) + C$ as an indivisible operation. This single functional unit reduces the latency for chained floating-point operations, as well as rounding errors and chip busing. It also reduces the number of adders/normalizers by combining the addition required for fast multiplication with accumulation. The MAF unit is made practical by a unique fast-shifter, which eases the overlap of multiplication and addition, and a leading-zero/one anticipator, which eases overlap of normalization and addition. The accumulate instruction required by this architecture reduces the instruction path length by combining two instructions into one. Additionally, the RS/6000 FPU is tightly coupled to the rest of the CPU, unlike typical floating-point coprocessor chips. As a result, floating-point and fixed-point instructions can be executed simultaneously. Load/store operations are performed using register renaming and store buffering to allow completely independent operation of load/store with arithmetic operations. Thus, data-cache accesses can occur in parallel with independent arithmetic operations. This unit attains a peak execution rate of 50 MFLOPS with a 25-MHz clock frequency and is capable of sustaining nearly that rate in complex programs such as graphics and Livermore loops.

Leading-zero anticipator (LZA) in the IBM RISC System/6000 floating-point execution unit *by E. Hokenek and R. K. Montoye, p. 71.* This paper presents a novel technique used in the multiply-add-fused (MAF) unit of the IBM RISC System/6000® (RS/6000) processor for normalizing the floating-point results. Unlike the conventional procedures applied thus far, the so-called leading-zero anticipator (LZA) of the RS/6000 carries out processing of the leading zeros and ones in parallel with floating-point addition. Therefore, the new circuitry reduces the total latency of the MAF unit by enabling the normalization and addition to take place in a single cycle.

Pseudorandom built-in self-test methodology and implementation for the IBM RISC System/6000 processor *by I. M. Ratiu and H. B. Bakoglu, p. 78.* This paper describes a unified self-test and system bring-up methodology. The components involved include a common on-chip processor (COP) that executes the chip self-test sequence and provides

an interface to the COP bus, a serial bus (COP bus) that links the chips to OCS and ESP, an on-card sequencer (OCS) that controls the self-test and system initialization sequences, and an engineering support processor (ESP) that is used for system verification, bring-up, and debug. Almost all RISC System/6000® chips contain embedded RAMs such as register files, caches, and directories; therefore, the self-test methodology described here is particularly suitable for logic chips that contain embedded arrays. Logic and RAM self-test is executed by a control processor (COP) integrated on the chips. The COP controls the self-test sequence, generates pseudorandom test vectors, scans them into chip registers, and provides the select lines that establish a one-to-one correspondence between RAM input/output and chip registers. The COP also drives RAM read/write lines during self-test, scans the captured RAM outputs, and compresses them to obtain a signature. After the vectors are scanned in, the chip runs for one or two system cycles, the logic outputs are captured in registers, and the chip state is scanned back into the COP, where it is compressed to obtain a signature. This procedure is repeated many times, and the final signature is then compared with a predetermined "good" signature to establish whether the chip is good or bad. Special techniques are developed to improve the coverage of logic that feeds RAMS or receives its inputs from RAMs. Both ac and dc self-test are described. The self-test sequence is controlled by a program stored in the OCS, and ESP is used during system bring-up to set up break-points and to display and modify the machine state.

Instruction scheduling for the IBM RISC System/6000 processor *by H. S. Warren, Jr., p. 85.* For fast execution on the IBM RISC System/6000® processor, instructions should be arranged in an order that uses the arithmetic units as efficiently as possible. This paper describes the scheduling requirements of the machine, and a scheduling algorithm for it that is used in two compilers.

Instruction scheduling beyond basic blocks *by M. C. Golumbic and V. Rainish, p. 93.* Instruction scheduling consists of the rearrangement or transformation of program statements, usually at the intermediate language or assembly code level, in order to reduce possible run-time delays between instructions. Such transformations must preserve data dependency and are subject to other constraints. Highly optimizing compilers employing instruction-scheduling techniques have proven to be effective in improving the performance of pipeline processors. Considerable attention has been given to scheduling code within the scope of basic blocks, i.e., straight-line sections of code. In this paper we present techniques for scheduling beyond basic blocks. This allows a further reduction in run-time delays such as those due, e.g., to branches and loops, enabling the exploiting of pipeline architectures which would not otherwise be possible.

Managing programs and libraries in AIX Version 3 for RISC System/6000 processors *by M. A. Auslander, p. 98.* This paper describes the program and program-library management facility that has been developed for the AIX® operating system, Version 3, as implemented for the IBM POWER (Performance Optimization With Enhanced RISC) architecture. It provides run-time loading of libraries, symbol resolution with type checking, and relocation. In addition, the use of the loader to add programs to an already running process or to the kernel is offered. The advantages of these functions and the techniques needed to provide a usable and efficient realization are described. Particular attention is given to the special problems posed by very large programs, and by very small programs which use services from very large libraries.

Evolution of storage facilities in AIX Version 3 for RISC System/6000 processors *by A. Chang, M. F. Mergen, R. K. Rader, J. A. Roberts, and S. L. Porter, p. 105.* The AIX® Version 3 storage facilities include features not found in other implementations of the UNIX® operating system. Maximum virtual memory is more than 1000 terabytes and is used pervasively to access all files and the meta-data of the file systems. Each separate file system (subtree) of the file name hierarchy occupies a logical disk volume, composed of space from possibly several disks. Database memory (a variant of virtual memory) and other database techniques are used to manage file system meta-data. These features provide the capacity to address large applications and many users, simplified program access to file data, efficient file buffering in memory, flexible management of disk space, and reliable file systems with short restart times.

Computation of elementary functions on the IBM RISC System/6000 processor *by P. W. Markstein, p. 111.* The additional speed and precision of the IBM RISC System/6000® floating-point unit have motivated reexamination of algorithms to perform division, square root, and the elementary functions. New results are obtained which avoid the necessity of doing special testing to get the last bit rounded correctly in accordance with all of the IEEE rounding modes in the case of division and square root. For the elementary function library, a technique is described for always getting the last bit rounded correctly in the selected IEEE rounding mode.

Volume 34, Number 2/3, 1990

Preface *by P. G. May, p. 139.*

Electro-optic sampling of high-speed devices and integrated circuits *by J. M. Wiesenfeld, p. 141.* The operating speeds of the fastest electronic devices and integrated circuits (ICs) have surpassed the capabilities of conventional electronic measurement instrumentation. Electro-optic sampling is an optical probing technique which has ultrashort temporal resolution and is capable of noninvasively probing ICs at internal nodes. This technique is voltage-sensitive because it relies upon the electric field produced by the signal

voltage on the device under test (DUT). The electric field (and hence the voltage) can be sampled because it produces birefringence in an electro-optic crystal which changes the state of polarization of an ultrashort-duration optical probe pulse that propagates through the electro-optic crystal. The electro-optic crystal is the substrate of the DUT for direct probing, is a crystal on a separate test structure for hybrid probing, and is a separate crystal placed above the DUT for external probing. Temporal resolution below 1 ps and a sensitivity below 0.1 mV/$\sqrt{\text{Hz}}$ have been demonstrated (though not in the same experiment). The principles of electro-optic sampling are reviewed in this paper. Selected applications for measurement of high-speed waveforms in discrete devices and in ICs are presented.

Picosecond noninvasive optical detection of internal electrical signals in flip-chip-mounted silicon integrated circuits *by H. K. Heinrich, p. 162.* This paper reviews the charge-sensing optical probing system, and shows how it may be used to detect internal current and voltage signals in flip-chip-mounted silicon integrated circuits. Previously, researchers have used this concept to detect both single-shot 200-MHz-bandwidth signals, without averaging, and 8-GHz-bandwidth stroboscopic signals. This system has a high sensitivity: 145-nA/$\sqrt{\text{Hz}}$ current sensitivity in typical bipolar transistors, and 1.35-mV/$\sqrt{\text{Hz}}$ voltage sensitivity in typical CMOS circuits (using a semiconductor laser probe). It is noninvasive, has a potential submicron spatial resolution, and should be capable of providing linear and calibrated measurements. Therefore, this probing approach should be a powerful tool for future circuit analysis and testing.

Picosecond photoemission probing of integrated circuits: Capabilities, limitations, and applications *by R. Clauberg, H. Beha, A. Blacha, and H. K. Seitz, p. 173.* The capabilities and limitations of the novel photoemission probing technique for signal measurements on internal nodes of VLSI integrated circuits are reviewed with respect to the range of possible applications of this method. Aspects such as voltage sensitivity, time resolution, minimum accessible feature size, sensitivity to perturbation effects, and impact on the circuit under test are considered. It is concluded that the especially high voltage sensitivity of this new method opens the field of diagnostics of circuits with ultrafast devices but partly low signal repetition rates, which is not accessible by other means. Such chips include, for example, complex logic chips and special telecommunication chips.

Flexible picosecond probing of integrated circuits with chopped electron beams *by D. Winkler, R. Schmitt, M. Brunner, and B. Lischke, p. 189.* The effective design and evaluation of high-speed integrated circuits is supported by internal noninvasive voltage-measurement techniques with picosecond time resolution. An electron-beam tester has therefore been developed which approaches the theoretical time-resolution limit of this method. It is based on the well-established e-beam technique for VLSI circuits, allowing for high flexibility in driving different kinds of high-frequency circuits under both conventional and critical conditions. The electron pulses of the stroboscopic test system are generated by a two-stage chopping system which was optimized to obtain very short pulses. It allows for a 7-ps effective pulse width which simultaneously yields a probe diameter of 0.5 μm and a probe current of 1 nA. This current results in a noise voltage of 20 mV when one period of a 1-GHz signal is recorded, with a total acquisition time of 0.1 s. Long-range phase shifting with high resolution is achieved by operating the upper stage of the blanking system at a high frequency and using the lower one as a selective gate. This allows propagation-delay measurements to be performed with a resolution of better than 2 ps over a range of several μs. The test system has thus far been used for that analysis of tunnel diodes, step-recovery diodes, bipolar frequency dividers, ring oscillators, and GaAs memories. Waveform measurement and evaluation at more than 60 different test points of a GaAs 1-kb SRAM in a six-hour session has demonstrated routine handling of complex high-speed circuit analysis.

Picosecond photoelectron microscope for high-speed testing of integrated circuits *by P. May, Y. Pastol, J.-M. Halbout, and G. Chiu, p. 204.* The performance of devices an circuits is advancing at a rapid pace with the advent of submicron design ground rules and switching times under 50 ps. The requirements for probing the internal nodes of these ultra-fast, -small, and -dense circuits give rise to great challenges for high-speed electron-beam testing. In this paper, we review the steps which have allowed electron-beam testing to achieve simultaneously 5-ps temporal resolution, 0.1-μm spot size, and 3 mV/$\sqrt{\text{Hz}}$ voltage sensitivity. The resulting newly developed instrument, called the picosecond photoelectron scanning electron microscope (PPSEM), is capable of measuring the state-of-the-art bipolar and FET circuits and also VLSI passive interconnects.

A submicron electron-beam tester for VLSI circuits beyond the 4-Mb DRAM *by F. Fox, J. Kölzer, J. Otto, and E. Plies, p. 215.* This paper describes the electron-optical low-voltage column of the submicron electron-beam tester. It can be used to produce an electron probe of 0.12-μm diameter, 2.5-nA probe current, and 1-kV beam voltage. It is shown that in the case of waveform measurements on 1.1-μm interconnection lines, the crosstalk is only ~ 3%. The voltage resolution is sufficient to allow the sense signal of a 4-Mb DRAM (dynamic random access memory) to be measured. Further internal measurements with the electron probe for the chip verification of the 4-Mb DRAM are also shown which demonstrate the flexibility and the benefits of electron-beam testing. On the basis of the measured performance data and its successful use in the circuit analysis of the 4-Mb DRAM, the submicron electron-beam tester appears to be suitable also for VLSI circuits with reduced design rules, e.g., for the 16-Mb DRAM. The improvements required for such future applications are briefly discussed.

Internal probing of submicron FETs and photoemission using individual oxide traps *by P. Restle and A. Gnudi, p. 227.* In submicron field-effect transistors with channel area less than 0.5 μm^2, the capture or emission of a single electron (or hole) in the gate oxide has an easily observable effect on the device resistance. Measurements are described in which the time and amplitude of the resistance change due to each capture and emission event from an individual trap are extracted to obtain the average capture and emission times, and the amplitude of the resistance change, at different temperatures, device biases, and light intensities. Techniques are described for using the data at different biases to characterize the trap, find the location of the trap in the device, and then use the trap as a probe of the oxide field (or surface potential) and the surface charge density within a 5–50-Å radius of the trap. In some devices a single trap can be resolved over almost all designed bias regions of the FET near room temperature. In effect, individual traps can be used as internal probes into VLSI devices of the present and future. Results from 2D computer device modeling of these devices are used to evaluate and understand these techniques. Methods for applying these techniques to the study of device degradation are discussed. Data are presented in which photoemission is observed from a single electron trap.

A submicron MOSFET parameter extraction technique *by B. El-Kareh, W. R. Tonti, and S. L. Titcomb, p. 243.* A technique is introduced for measuring electron and hole mobilities as a function of temperature and normal field in inverted silicon surfaces. We also introduce a new definition of threshold voltage which allows the method to measure mobility independent of channel dimensions and resistance in series with the channel. The results are used to extract the resistance in series with the channel, the effective channel dimensions, and the intrinsic MOSFET transconductance. The technique is demonstrated on MOSFETs with channel lengths ranging from 0.25 μm to 20 μm.

Electron-beam technology for open/short testing of multi-chip substrates *by S. D. Golladay, N. A. Wagner, J. R. Rudert, and R. N. Schmidt, p. 250.* We discuss the need for noncontact electrical testing of high-performance multi-chip substrates and describe an electron-beam tester developed for this application. We describe the operational principles of the tester and compare and contrast its performance with that of mechanical probe testers. Finally, we discuss the motivations and technical issues involved in extending the electron-beam test method to future high-performance packages.

The development of ultra-high-frequency VLSI device test systems *by C. W. Rodriguez and D. E. Hoffman, p. 260.* The development of test systems for high-performance semiconductor logic and memory devices is discussed. The capabilities of shared-resource and tester-per-pin system architectures are reviewed. Test-system hardware design to provide high-speed pin electronics and generation of LSSD, weighed random, and algorithmic patterns is described. The reasons for the selection of the tester-per-pin system architecture are given in terms of the way in which overall system accuracy and test-system user flexibility are maximized for differing test methodologies.

Test generation for VLSI chips with embedded memories *by E. K. Vida-Torku, J. A. Monzel, T. L. Bossis, C. E. Radke, and D. M. Wu, p. 276.* An effective approach for generating patterns for testing memories embedded in logic is presented. Through circuit testability analysis, which is a study of the effect of process defects on memory circuits, unique algorithms can be derived for testing the memory. Circuit and logic designs for test features that are required to make the pattern generation process optimal are discussed. An analytical method is described which assesses the performance characteristics of the memory after functional test.

Simulation of embedded memories by defective hashing *by L. M. Huisman, p. 289.* Because logic designs are becoming more complex and extensive, they increasingly tend to contain embedded memories. In the simulation (particularly fault simulation) of these designs, the embedded memories may be found to require large amounts of storage unless a carefully designed simulation strategy is adopted. This paper describes a technique that drastically reduces the storage required in the fault simulation of such large designs. The required amount of storage can be fixed at compile time or at load time, and can almost always be made to fit in the available storage at the cost of only a small decrease in the predicted exposure probabilities.

A logic chip delay-test method based on system timing *by F. Motika, N. N. Tendolkar, C. C. Beh, W. R. Heller, C. E. Radke, and P. J. Nigh, p. 299.* In this paper we present a novel approach to delay-testing of VLSI logic chips based on the level-sensitive scan design (LSSD) methodology. The objective of the delay test is to reduce significantly the failures of multi-chip modules at system integration test while minimizing the complexity and cost of subassembly testing. Because system timing data are used to derive test specifications, the delay defects that are most likely to cause a system path failure are detected a high percentage of the time. With the implementation of the delay test in the wafer production line, the system final-test failure rate of multi-chip modules used in IBM mainframe machines has dropped significantly.

An ac test structure for fast memory arrays *by R. C. Wong, p. 314.* An ac test structure (ACTS) built into fast memory arrays is required to make them truly ac-testable, with 5–10% timing accuracy. Since their ac performance is very difficult to characterize, wafer tester timing uncertainty is generally about 10–50% of a typical array access time. More accurate testers are complex and expensive; they require long development time and have complicated operation procedures. ACTS is a simpler, cheaper, and more practical means of achieving greater accuracy. An ac test structure is composed

of a tunable timer and path-shifting oscillators (PSOs) built around the various access paths of the array. The timer generates the array clocks with adjustable pulse widths, and the PSOs transform time intervals into frequencies. In the future, tester accuracy will improve, but memory performance will have accelerated even more. Thus, the need for ACTS is critical and will remain so in the foreseeable future.

Gross delay defect evaluation for a CMOS logic design system product *by O. Bula, J. Moser, J. Trinko, M. Weissman, and F. Woytowich, p. 325.* Randomly occurring gross delay defects allow chips to pass full stuck-fault testing at both wafer and module levels, but cause them to fail when operated at system speeds. This paper describes the results of an experiment designed to determine the actual delay defect component of shipped product quality level (SPQL) for a CMOS combination standard cell/gate array design system. More than 60,000 modules, representing chips from the same IBM computer system, have been delay-tested using the technique presented in this paper. The test technique uses the stuck-fault patterns for level-sensitive scan design (LSSD) product. The stuck-fault patterns are modified or "twisted" according to specific algorithms to propagate transitions through paths just prior to the output measure. The patterns are applied at system speed timings provided by the chip designers. Any gross delay defect present in a tested path causes a fail. The failing modules were characterized to determine the size of the delay defects. Failure diagnostics were performed on the defective modules by using existing stuck-fault diagnostic tools and a development version of a transition fault simulator. These were sent to physical failure analysis for delayering, visual verification, and electrical characterization. A summary of physical defects which produced gross delay defects is presented.

Boundary-scan design principles for efficient LSSD ASIC testing *by R. W. Bassett, M. E. Turner, J. H. Panner, P. S. Gillis, S. F. Oakland, and D. W. Stout, p. 339.* A boundary-scan logic design method that depends only on level-sensitive scan design (LSSD) principles has been developed for IBM CMOS application-specific integrated circuit (ASIC) products. This technique permits comprehensive testing of LSSD ASICs with high signal input/output (I/O) pin counts, using relatively inexpensive reduced-pin-count automatic test equipment (ATE). This paper describes the LSSD logic structures required, the reduced-pin-count testing and burn-in processes used, and the ASIC product design decisions that must be made to establish a consistent boundary-scan implementation.

Design for testability and diagnosis in a VLSI CMOS System/370 processor *by C. W. Starke, p. 355.* This paper describes the design for testability and diagnosis in an IBM System/370 processor based on VLSI CMOS technology. The design incorporates built-in pseudorandom-pattern self-test and the boundary-scan technique. This technique permits the migration of tests generated for component-level to higher-level packages such as printed circuit boards and the system. Consequently, the expense for testing of higher-level packages can be reduced, and the test equipment for the processor can be simplified. The design also offers economical diagnostic capability.

Aliasing errors in linear automata used as multiple-input signature analyzers *by W. Daehn, T. W. Williams, and K. D. Wagner, p. 363.* This paper deals with the aliasing probability of multiple-input data compressors used in self-testing networks. It is shown that a far more general class of linear machines than linear feedback shift registers can be used for data compression purposes. The function of these machines is modeled by a Markov process. The steady-state value of the aliasing probability is shown to be the same as for single-input signature analysis registers. An easily verifiable criterion is given that allows one to decide whether a given linear machine falls into this class of multiple-input data compressors. The steady-state value of the aliasing probability is shown to be independent of the correlation of the data streams at the inputs of the data compressor. Two kinds of circuits are analyzed in more detail with respect to their aliasing properties: linear feedback shift registers with multiple inputs, and linear cellular automata. Simulation results show the effect of the next-state function on the steady-state value of the aliasing probability and the effect of correlation on the transient response.

Improved cutting algorithm *by J. Savir, p. 381.* The cutting algorithm allows computation of bounds on signal probabilities and detection probabilities in combinational networks. These bounds can be used to determine the necessary pseudorandom test length needed to test a network. One of the problems with the cutting algorithm is that it may compute loose bounds which translate into unnecessarily long test lengths. The object of this paper is to improve the cutting algorithm so that the computed bounds become satisfactory. The improved cutting algorithm is a careful combination of the original cutting algorithm and the Parker–McCluskey algorithm. The tightness of the computed bounds may vary depending on which portion of the circuit is handled with the cutting algorithm and which with the Parker–McCluskey algorithm. Thus, the user of the improved cutting algorithm can actually control and trade off the accuracy of the results against the computational effort needed to achieve them.

Cellular automata circuits for built-in self-test *by P. D. Hortensius, R. D. McLeod, and B. W. Podaima, p. 389.* Results are presented for a variation on a built-in self-test (BIST) technique based upon a distributed pseudorandom number generator derived from a one-dimensional cellular automata (CA) array. These cellular automata logic block observation (CALBO) circuits provide an alternative to conventional design for testability circuitry such as built-in logic block observation (BILBO) as a direct consequence of reduced cross-correlation between the bit streams which are used as inputs to the logic unit under test. The issue of

generating probabilistically weighted test patterns for use in built-in self test is also addressed. The methodology presented considers the suitability of incorporating structures based on cellular automata, a strategy which, in general, improves test pattern quality. Thus, CA-based structures quality as attractive candidates for use in weighted test pattern generator design. The analysis involved in determining and statistically evaluating these potential models is discussed, and is compared with that for previous as well as statistically independent models. Relevant signature analysis properties for elementary one-dimensional cellular automata are also discussed. It is found that cellular automata with cyclic-group rules provide signature analysis properties comparable to those of the linear feedback shift register. The results presented here are based upon simulation.

Built-in self-test support in the IBM Engineering Design System *by B. L. Keller and T. J. Snethen, p. 406.* To evaluate the effectiveness of built-in self-test (BIST) for logic circuits, the test design automation (TDA) group within the IBM Engineering Design System (EDS) has developed tools to support BIST. This paper is an overview of that support. The specific hardware approaches taken are described briefly, and a short description is given of the major tools that have been developed and the methodology for using them. The performance of the system is shown for two sample circuits.

Self-testing the 16-Mbps adapter chip for the IBM token-ring local area network *by S. F. Oakland, J. L. Corr, J. D. Blair, V. R. Norman, and W. J. DeGuise, p. 416.* This paper describes the boundary-scan and built-in self-test (BIST) functions of the IBM token-ring local area network (LAN) adapter chip. These functions present a number of unique features. First, less that 1% of available standard cell circuits were needed to implement these functions. Second, clocking methods used in different logical macros were merged into a comprehensive clocking sequence for self-test. Finally, asynchronous serial and parallel interfaces were provided to facilitate the communication between a test system and the chip's built-in test circuits. Although self-test and boundary-scan provide for an inexpensive higher-level package test, evaluation showed that automatically generated deterministic patterns provide a better-quality VLSI chip manufacturing test.

LAN interface chip and mixed-signal testing developments *by J. J. Van Horn, R. A. Waller, R. J. Prilik, and K. C. Bocash, p. 428.* This paper describes the local area network (LAN) semiconductor chip, from early development through volume manufacturing production, highlighting the testing achievements associated with its qualification and release to production. The paper is divided into two sections dealing with laboratory development and production in manufacturing. The development section describes the importance of partitioning via latch-based boundaries, and how it reduced development cycle time by allowing independent debug and diagnostic tests. The development of tests to characterize the phase-lock loop (PLL) at application conditions and the evolution of these tests into efficient production test vehicles are discussed. Techniques are described that provide new approaches to analog testing, focusing on adapting a production digital tester to meet the characterization and production requirements of a very sensitive PLL/receiver/transmitter analog design. The manufacturing discussion is centered around three areas, wafer testing, yield learning, and reliability. Unique concepts in each area are presented, along with detailed descriptions of specific applications.

Volume 34, Number 4, 1990

Preface *by M. Heiblum and S. S. Iyer, p. 450.*

Experimental technology and performance of 0.1-μm-gate-length FETs operated at liquid-nitrogen temperature *by G. A. Sai-Halasz, M. R. Wordeman, D. P. Kern, S. A. Rishton, E. Ganin, T. H. P. Chang, and R. H. Dennard, p. 452.* An overview is presented of our work to explore the extendibility of the silicon FET technology to the 0.1-μm-gate-length level. Self-aligned, *n*-channel, polysilicon-gated FETs were designed for operation at 77K, with reduced power-supply voltage. Direct-write electron-beam lithography was used to pattern all levels, while other processing followed established lines. Noteworthy results of the work included the observation of a clear manifestation of velocity overshoot, which contributed to achieving extrinsic transconductances above 940 μS/μm at 0.07-μm gate length. The measured switching delay of ring oscillators which contained 0.1-μm-gate-length devices was as low as 13.1 ps, with simulations showing potential for reduction to below 5 ps. Both the transconductance and the switching times are the best values observed for FETs to date—indicating continuing value in the scaling of FETs to dimensions well beyond those currently used.

Monte Carlo analysis of semiconductor devices: The DAMOCLES program *by S. E. Laux, M. V. Fischetti, and D. J. Frank, p. 466.* The behavior of small semiconductor devices is simulated using an advanced Monte Carlo carrier transport model. The model improves upon the state of the art by including the full band structure of the semiconductor, by using scattering rates computed consistently with the band structure, and by accounting for both long- and short-range interactions between carriers. It is sufficiently flexible to describe both unipolar and bipolar device operation, for a variety of semiconductor materials and device structures. Various results obtained with the associated DAMOCLES program for *n*- and *p*-channel Si MOSFETs, GaAs MESFETs, and Si bipolar junction transistors are presented.

Submicron-gate-length GaAs MESFETs *by T. N. Jackson, B. J. Van Zeghbroeck, G. Pepper, J. F. DeGelormo, T. Keuch, H. Meier, and P. Wolf, p. 495.* It is well known that reducing gate length is a powerful means to increase the

transconductance and transit frequency of GaAs MESFET devices. However, by reducing the gate length without scaling channel doping and thickness, the performance obtained is limited by short-channel effects and parasitics. In this paper we present an overview of our work on two different MESFET structures, illustrating how device performance can be increased by decreasing the gate length, with the result that appropriately scaled MESFETs compare favorably with GaAs–AlGaAs heterojunction FETs. From our work—including some recent results on 0.15-μm-gate-length implantation-self-aligned MESFETs—we conclude that it should be possible to increase the speed of high-speed GaAs MESFET (logic, analog, and microwave) circuits through the use of devices having gate lengths less than 0.5 μm.

Heterojunction FETs in III–V compounds *by R. A. Kiehl, P. M. Solomon, and D. J. Frank, p. 506.* We review work on heterojunction FETs (HFETs) fabricated from III–V compounds, with emphasis on the unique properties of such devices and their applicability to high-speed logic circuits. After discussing their general properties, including their uniquely high carrier mobility and fast switching speed, we discuss HFETs investigated at the IBM Thomas J. Watson Research Center, i.e., the semiconductor–insulator–semiconductor FET (SISFET) and quantum-well metal–insulator–semiconductor FET (QW-MISFET)—and their possible circuit applications. Finally, the opportunities for achieving a circuit performance level beyond that offered by the GaAs–(Al,Ga)As materials system are explored.

Ballistic hot-electron transistors *by M. Heiblum and M. V. Fischetti, p. 530.* We present an overview of work at the IBM Thomas J. Watson Research Center on the tunneling hot-electron transfer amplifier (THETA) device—including its use as an amplifier and as a tool for investigating ballistic hot-electron transport. In the initial, vertically configured version of the device, a quasi-monoenergetic, variable-energy, hot-electron beam is generated (via tunneling) which traverses a thin GaAs region and is then collected and energy-analyzed. As the hot electrons traverse the device, they are used to probe scattering events, band nonparabolicity, size-quantization effects, and intervalley transfer. A recent, lateral version of the device has been used to demonstrate the existence of ballistic hot-electron transport in the plane of a two-dimensional electron gas, and the associated possibility of achieving high gain.

Compound semiconductor heterostructure bipolar transistors *by S. Tiwari, S. L. Wright, and D. J. Frank, p. 550.* This paper is primarily an overview of our work on the technology, material and electronic properties, and performance limitations of compound semiconductor heterostructure bipolar transistors. Graded-gap epitaxial *n*-type ohmic contacts and *p*-type shallow diffusion ohmic contacts are important in the fabrication of high-performance (Al,Ga)As/GaAs devices. In the device structure implemented, the presence of a wide-gap *p*-type (Al,Ga)As extrinsic base region at the surface suppresses surface recombination, thereby enhancing the current gain at small device dimensions. We discuss experimental and theoretical results concerning the limiting physical effects due to heterostructure design and intrinsic and extrinsic bulk phenomena of compound semiconductors, emphasizing the understanding developed and the discoveries made during the course of our efforts. As device speeds have increased with coordinated scaling, dispersive effects have become increasingly important. We show how these may be included by modifying the conventional quasi-static modeling of the bipolar transistor, in order to obtain a realistic simulation of fast switching transients. Finally, we discuss scaling of heterostructure bipolar transistors, and implications of the use of lower-bandgap materials and operation at cryogenic temperatures.

High-speed GaAs/AlGaAs optoelectronic devices for computer applications *by C. S. Harder, B. J. Van Zeghbroeck, M. P. Kesler, H. P. Meier, P. Vettiger, D. J. Webb, and P. Wolf, p. 568.* We present an overview, mainly of work in our laboratory, of low-threshold GaAs/AlGaAs quantum-well laser diodes and GaAs metal–semiconductor–metal photodetectors—two optoelectronic devices which show good promise for use in computer-related communication. Present-day telecommunication device technology (based on InP materials) is not well suited to the requirements of optical data communication among and within computers because the computer environment is much more demanding. It imposes a higher ambient temperature on the devices, and requires denser packaging and smaller power dissipation per device, as well as a high degree of parallelism. The GaAs/AlGaAs device technology is ideally suited to this task because of the possibility of integration of arrays of high-speed, low-threshold laser diodes and high-speed photodetectors with high-performance electronic circuits.

An electromagnetic approach for modeling high-performance computer packages *by B. J. Rubin, p. 585.* Described here is an electromagnetic approach for the analysis of high-performance computer packages such as the thermal conduction module (TCM) used in the IBM 3080 and 3090 processor units. Modeling of signal paths and limitations of previous methods are discussed. Numerical results are presented for propagation characteristics associated with signal lines and vias, and for coupled noise between signal lines. The results are compared with those obtained by means of test vehicles, scale models, and capacitance calculations.

High-speed signal propagation on lossy transmission lines *by A. Deutsch, G. V. Kopcsay, V. A. Ranieri, J. K. Cataldo, E. A. Galligan, W. S. Graham, R. P. McGouey, S. L. Nunes, J. R. Paraszczak, J. J. Ritsko, R. J. Serino, D. Y. Shih, and J. S. Wilczynski, p. 601.* This paper addresses some of the problems

encountered in propagating high-speed signals on lossy transmission lines encountered in high-performance computers. A technique is described for including frequency-dependent losses, such as skin effect and dielectric dispersion, in transmission line analyses. The disjoint group of available tools is brought together, and their relevance to the propagation of high-speed pulses in digital circuit applications is explained. Guidelines are given for different interconnection technologies to indicate where the onset of severe dispersion takes place. Experimental structures have been built and tested, and this paper reports on their electrical performance and demonstrates the agreement between measured data and waveforms derived from analysis. The paper addresses the problems found on lossy lines, such as reflections, rise-time slowdown, increased delay, attenuation, and crosstalk, and suggests methods for controlling these effects in order to maintain distortion-free propagation of high-speed signals.

Volume 34, Number 5, 1990

Preface *by M. A. Wesley, p. 634.*

A modeling system for top-down design of assembled products *by M. Mäntylä, p. 636.* The design of a mechanical product usually takes place primarily in a top-down fashion, where the designer first generates a rough, overall sketch of the product and its main components. Later, the designer refines the sketch to a detailed level while taking into account the relevant requirements posed by strength, cost, manufacturability, serviceability, and other similar considerations. Current computer-aided design (CAD) systems provide only limited support for this kind of work. For instance, they cannot deal with geometric or other information at varying levels of detail, nor do they capture explicitly geometric relationships among components intended to be joined together in an assembly. This paper describes early results of ongoing research on supporting top-down design of mechanical products and discusses the major requirements for CAD systems used for top-down design. A prototype design system is described that provides the following characteristics not usually found in ordinary CAD systems: structuring of product information in several layers, according to the stage of the design process; representation of geometric information about components at several levels of detail; and representation and maintenance of geometric relationships of components by means of a constraint-satisfaction mechanism.

Computer-aided design of slider bearings in magnetic disk files *by K. L. Deckert, p. 660.* This paper reviews the application of lubrication theory to slider bearings in magnetic disk files. For more than thirty years, slider bearings have been used to maintain close and precise spacings between recording transducer and recording medium in disk files. Computer modeling has been central to the design and performance analysis of these systems. The topics covered are the basic design, sensitivity and tolerance analysis, dynamic characteristics, and response to disk excitations from the disk. The main purpose of this paper is to review the use of computer modeling in design of slider bearings; however, the discussion of slider modes in the slider dynamics section is new.

Disk file access-time constraints imposed by magnetic air-bearing compliance *by E. S. Cooper, p. 668.* Use is commonly made of coil voltages which produce maximum acceleration and deceleration to wrest the fastest access performance from the actuator of a disk file. Alternatively, equivalent performance is easily obtained with less control optimality by simply increasing the coil voltage. A limit to greater coil voltage, however, arises from the need to avoid harmful effects when the servo loses control of the actuator, the actuator slams into its crash stop, and the crash force stresses the magnetic head air bearing. Magnetic head air-bearing compliance, therefore, is one of three fundamental constraints that limit the access-time performance of an actuator. To obtain improved access-time performance from a disk file, the following should be optimized: the air-bearing slider, the actuator crash stop, and the actuator mechanical time constant. This paper presents relevant design considerations.

A simple finite element model for reactive sputter-deposition systems *by F. Jones and J. S. Logan, p. 680.* The method of finite element analysis is used to calculate oxygen-concentration profiles and oxygen-absorption rates at the substrate and along the shields in an rf magnetron reactive sputtering system having a 12-inch-diameter magnetron. Results for several shield arrangements are calculated as a function of oxygen flow rate. The approach used assumes the following: (i) The target is being sputtered in the metallic state; (ii) the oxygen-concentration profile in the system can be calculated from the diffusion equation; (iii) the maximum amount of oxygen that can be absorbed at any point in the system is proportional to the metal deposition rate at that point; (iv) the target absorbs no oxygen as long as it is in the metallic state. The relative metal deposition rate along the substrate and shields is calculated, normalized to the measured deposition rate at the substrate, and used as a boundary condition for the diffusion equation. The calculated oxygen flow rate for the formation of stoichiometric substrate films agrees with experimental results to within 15%. The critical flow rate at which the target oxidizes, Q_c, is measured experimentally and when used in the model gives an oxygen partial pressure of about $0.31 \cdot 10^{-6}$ torr at the sputtering track.

Computation of current distribution in electrodeposition, a review *by J. O. Dukovic, p. 693.* This paper reviews the research literature that has appeared since 1980 on computer calculations of current distribution in the field of electrodeposition. Key contributions and general trends are identified, with particular emphasis given to applications in

the electronics industry. The survey reveals how numerical models have provided the technology of electrodeposition with general understanding, predictive power, and the capability of optimizing deposit uniformity. Anticipated developments for the nineties are discussed.

Finite element analysis of planar stress anisotropy and thermal behavior in thin films *by K. F. Young, p. 706.* To show the capability and diversity of finite element analysis, we calculate three-dimensional planar anisotropic stress distributions for various thin-film geometries and materials in response to thermal and electrical stimuli, for specific boundary conditions. The simulated residual film stresses are verified with acoustic microscopy measurements, substrate flexure measurements, and the use of thermal environment techniques. Simple shapes are analyzed as building blocks for more complex structures. Effects of nonlinear electrical resistance are also analyzed.

Thermoelastic behavior of X-ray lithography masks during irradiation *by I. A. Shareef, J. R. Maldonado, Y. Vladimirsky, and D. L. Katcoff, p. 718.* This paper presents computer calculations of thermoelastic effects in X-ray lithography masks caused by the absorption of X-rays during exposure. Several mask structures are considered, with different substrate and absorber materials, using finite element analysis. Part I of the paper deals with short-pulse X-ray irradiation (e.g., from gas plasma, laser-heated plasma, or exploding wire sources), and Part II describes irradiation during exposure with a synchrotron-storage-ring X-ray source. For the short-pulse irradiation, results indicate a maximum rise in temperature on the mask of about 30°C for a 2-ns exposure with a 10-mJ/cm^2 X-ray pulse. Mechanical static analysis shows that the maximum stress in the absorber films, which is due to maximum temperature differences in the mask layers, occurs at the end of the pulse. The magnitude of the induced thermoelastic stress is found comparable to the intrinsic stress level of the mask materials (typically 2–5 $\times$ 10^8 dyn/cm^2). The analysis indicates that when pulse amplitudes reach 10 mJ/cm^2, there will be a need for experimental study of X-ray mask distortion during exposure to short X-ray pulses. A one-dimensional model is developed for the case of storage-ring irradiation. The model predicts distortions of the printed image due to a thermal wave developed on the mask by scanning on the X-ray beam. Experimental results are presented showing that the effect is negligible under normal operating conditions but may become noticeable for operation in vacuum or without proper heat sinks.

Pythagorean hodographs *by R. T. Farouki and T. Sakkalis, p. 736.* The hodograph of a plane parametric curve $\mathbf{r}(t) = \{x(t), y(t)\}$ is the locus described by the first parametric derivative $\mathbf{r}'(t) = \{x'(t), y'(t)\}$ of that curve. A *polynomial* parametric curve is said to have a *Pythagorean* hodograph if there exists a polynomial $\sigma(t)$ such that $x'^2(t) + y'^2(t) \equiv \sigma^2(t)$, i.e., $(x'(t), y'(t), \sigma(t))$ form a "Pythagorean triple." Although Pythagorean-hodograph curves have fewer degrees of freedom than general polynomial curves of the same degree, they exhibit remarkable attractive properties for practical use. For example, their arc length is expressible as a *polynomial* function of the parameter, and their offsets are *rational* curves. We present a sufficient-and-necessary algebraic characterization of the Pythagorean-hodograph property, analyze its geometric implications in terms of Bernstein–Bézier forms, and survey the useful attributes it entails in various applications.

Finding compact coordinate representations for polygons and polyhedra *by V. J. Milenkovic and L. R. Nackman, p. 753.* Practical solid modeling systems are plagued by numerical problems that arise from using floating-point arithmetic. For example, polyhedral solids are often represented by a combination of geometric and combinatorial information. The geometric information may consist of explicit plane equations, with floating-point coefficients; the combinatorial information may consist of face, edge, and vertex adjacencies and orientations, with edges defined by face–face adjacencies and vertices by edge–edge adjacencies. Problems arise when numerical roundoff error in geometric operations causes the geometric information to become inconsistent with the combinatorial information. These problems can be avoided by using exact arithmetic instead of floating-point arithmetic. However, some operations, such as rotation, increase the number of bits required to represent the plane equation coefficients. Since the execution time of exact arithmetic operators increases with the number of bits in the operands, the increased number of bits in the plane equation coefficients can cause performance problems. One proposed solution to this performance problem is to *round* the plane equation coefficients without altering the combinatorial information. We show that such rounding is NP-complete.

Finding the distance between two circles in three-dimensional space *by C. A. Neff, p. 770.* In this paper we investigate, from an algebraic point of view, the problem of finding the distance between two circles located in *double*R^3. We show, by combining a theorem about solvable permutation groups and some explicit calculations with a computer algebra system, that, in general, the distance between two circles is an algebraic function of the parameters defining them, but that this function is *not* solvable in terms of radicals. Although this result implies that one cannot find a "closed-form" solution for the distance between an arbitrary pair of circles in *double*R^3, we discuss how such an algebraic quantity can still be manipulated symbolically by combining standard polynomial operations with an algorithm for isolating the real roots of a polynomial in a convenient data structure for real algebraic numbers. This data structure and its operations have been implemented.

Volume 34, Number 6, 1990

Preface *by T. F. Kuech, p. 794.*

Design of low-temperature thermal chemical vapor deposition processes *by D. B. Beach, p. 795.* The importance of an integrated approach involving synthetic chemistry, physical chemistry, and chemical engineering to the development of new thermal chemical vapor deposition (CVD) processes for the production of thin-film electronic materials is discussed. Particular emphasis is placed on choosing precursor molecules with facile thermal decomposition pathways that lead to pure films at low temperatures. Two examples from our laboratories, the deposition of copper from trialkylphosphine cyclopentadienyl copper complexes and the deposition of gallium nitride from diethylgallium azide, are used to illustrate the principles of precursor selection, the design factors for the construction of thermal CVD reactors, and the selection of processing conditions that optimize production of the desired material. In addition, new work on the thermal CVD of copper using an advanced reactor and examples of selective copper deposition and conformal copper deposition using the reactor are presented.

Low-temperature Si and Si:Ge epitaxy by ultrahigh-vacuum/chemical vapor deposition: Process fundamentals *by B. S. Meyerson, p. 806.* This paper is an overview of work at the IBM Thomas J. Watson Research Center on the chemical and physical considerations underlying the development of a low-temperature chemical vapor deposition process, designated ultra high-vacuum/chemical vapor deposition (UHV/CVD). The origins of the rigorous vacuum and chemical purity requirements of the process are discussed. Operating in the range of 500°C, the process has made it possible to explore the use, in silicon-based devices and atomic-length-scale structures, of a number of metastable materials in the Si:Ge:B system. Also discussed is associated experimental work on the fabrication of high-speed heterojunction bipolar transistors and high-mobility two-dimensional hole-gas structures.

Selective epitaxial growth of silicon and some potential applications *by B. J. Ginsberg, J. Burghartz, G. B. Bronner, and S. R. Mader, p. 816.* In the selective epitaxial growth (SEG) of silicon, growth occurs only on exposed silicon areas of a silicon substrate. Substrate regions on which silicon growth is not desired are masked by a dielectric film, typically silicon dioxide or silicon nitride. Use of the process permits the fabrication of novel silicon devices and integrated-circuit structures. In this paper, an overview is presented of our studies on the SEG process at the IBM Thomas J. Watson Research Center. Aspects covered are the kinetics of the process using a $SiCl_4$ and H_2 gaseous mixture, the associated suppression of deposition on silicon dioxide and silicon nitride, and some potential applications of the process to the fabrication of bipolar devices and dynamic random access memory (DRAM) cells.

Advances in metalorganic vapor-phase epitaxy *by M. A. Tischler, p. 828.* Metalorganic vapor-phase epitaxy (MOVPE) has become a well-established technique for the epitaxial growth of layers of III–V compound semiconductors since its introduction in 1968. Use has been made of the technique to produce such layers and associated devices to very demanding specifications. This paper describes MOVPE, followed by an overview of work at the IBM Thomas J. Watson Research Center on the technique, with emphasis on doping and selective epitaxy. Device applications are included to highlight the need to take into account the influence of materials and growth phenomena in order to produce optimum devices.

Diamondlike carbon films by rf plasma-assisted chemical vapor deposition from acetylene *by A. Grill, B. S. Meyerson, and V. V. Patel, p. 849.* This paper is an overview of studies performed at the IBM Thomas J. Watson Research Center on diamondlike carbon (hydrogenated amorphous carbon) films, including some recent results on their tribological properties. The films were prepared by rf plasma-assisted chemical vapor deposition (PACVD) from acetylene. Their structure and composition were characterized by a variety of methods such as X-ray and TEM diffractometry, XPS, high-resolution ^{13}C NMR spectroscopy, and $H(^{15}N\alpha, \gamma)C$ nuclear-reaction profiling. Their adhesion to various substrates, coefficients of static and dynamic friction, and wear resistance were also characterized. It was found that the films were essentially amorphous, reaching their bulk composition after 40 nm of growth above the initial growth interface. Their bulk composition included about 40% hydrogen. More diamondlike carbon bonding was obtained at the initial growth interface than in the bulk range. The ratio of sp^2 to sp^3 carbon atoms was found to be 1.6, with virtually all sp^3 carbon atoms bound to one or more hydrogen atoms. The diamondlike carbon films (DLC) displayed excellent adhesion to the surface of Si. Furthermore, the films could be bonded to films of otherwise nonbonding metals, provided the metals formed stable silicides. By using an interfacial Si film several atomic layers thick, adhesion to the metal films could be improved to the extent that attempts to detach the DLC films fractured the underlying metal films. The adhesive DLC films were found to have a very high resistance to sliding wear and chemical attack, and are therefore useful in various applications as very thin protective coatings.

On-chip wiring for VLSI: Status and directions *by M. B. Small and D. J. Pearson, p. 858.* The thirty-year history of silicon integrated circuits has resulted in dramatic increases in both the number of devices per chip and circuit speed. A consequence of scaling to submicron dimensions is that the major component of propagation delay will transfer from the devices to the interconnecting "wires." Additionally, increased integration, together with scaling, leads to a need for more numerous interconnections on a chip and higher current densities. Accommodation to these changes will necessitate the use of new materials arranged in three-dimensional wiring structures which have the ability to make the most effective use of the area of the chip. Generic processing routes to

achieve the desired structures are reviewed and examples are presented of two experimental structures with layers of planar wiring and vertical vias between planes. One of these integrates aluminum-alloy wiring with tungsten vias in a silicon dioxide dielectric; the other integrates copper wiring and vias in polyimide dielectric with the goal of minimizing delay due to on-chip wiring.

Surface and interfacial energies of $CoSi_2$ and Si films: Implications regarding formation of three-dimensional silicon-silicide structures *by K.-N. Tu, p. 868.* Formation of three-dimensional, multilevel structures consisting of epitaxial silicon and silicide films is currently of interest in the microelectronics technology. However, such structures have been difficult to produce because of surface wetting differences. To obtain associated surface energy information, an analysis was carried out of published data on the kinetics of crystallization of amorphous $CoSi_2$ and Si films. The analysis indicated that the amorphous-to-crystalline interfacial energy of amorphous $CoSi_2$ films is about one-fourth that of amorphous Si films, from which it was inferred that the surface energy of epitaxial $CoSi_2$ films is less than that of epitaxial Si films. The approach used in the analysis is general and should be extendable to other systems.

Surface chemistry of the WF_6-based chemical vapor deposition of tungsten *by M. L. Yu, K. Y. Ahn, and R. V. Joshi, p. 875.* This paper is an overview of work at the IBM Thomas J. Watson Research Center on chemical aspects of the WF_6-based chemical vapor deposition of tungsten. The focus is on two deposition processes. In the first process, tungsten deposition occurs through the chemical reduction of WF_6 by a silicon substrate. The thickness of the tungsten film thus grown is limited by the transport of silicon through the deposited film. In the second process, deposition occurs in a WF_6–silane mixture by the following reactions: reduction of adsorbed WF_6 by a surface layer of silicon to form tungsten, and the concurrent dehydrogenation of the adsorbed silane to replenish the silicon. This process permits the deposition of tungsten on any substrate, provided the initial reaction of the substrate with the WF_6–silane mixture can provide a tungsten or silicon "seed" layer to initiate the reaction cycle. Furthermore, the process is not limited by tungsten film thickness and hence permits the deposition of relatively thick films. Although surface selectivity is possible with regard to materials such as high-quality SiO_2, on which such a seed layer cannot be formed, the selectivity is adversely affected by the presence of particulates.

Magnetic thin films in recording technology *by V. S. Speriosu, D. A. Herman, Jr., I. L. Sanders, and T. Yogi, p. 884.* This paper is a review of recent progress in magnetic thin films for use in recording media and heads. Emphasis is on work that has been carried out at IBM. Topics covered include thin-film media for high-density recording, laminated soft-magnetic films for controlling domains and extending the frequency range of inductive heads, exchange-biasing of magnetoresistive sensors, and magnetic multilayer structures.

Magnetic multilayer structures *by R. F. C. Farrow, C. H. Lee, and S. S. P. Parkin, p. 903.* This is an overview of work at the IBM Almaden Research Center on magnetic multilayer structures comprising one, several, or many magnetic films sandwiched between nonmagnetic films. In recent years there has been increasing interest in such structures because of their novel and potentially useful properties. Recent examples of magnetic multilayer structures grown by molecular beam epitaxy (MBE) and sputtering are described. It is seen that MBE and sputtering are complementary techniques for the preparation of such structures, and that the ability to modify their magnetic properties by suitably designing their architecture is a key to their further development.

Lanthanide gallate perovskite-type substrates for epitaxial, high-T_c superconducting $Ba_2YCu_3O_{7-\delta}$ films *by E. A. Giess, R. L. Sandstrom, W. J. Gallagher, A. Gupta, S. L. Shinde, R. F. Cook, E. I. Cooper, E. J. M. O'Sullivan, J. M. Roldan, A. P. Segmüller, and J. Angilello, p. 916.* Previous studies had indicated promising use of lanthanide gallate perovskite-type substrates for the deposition of epitaxial, high-T_c superconducting $Ba_2YCu_3O_{7-\delta}$ (BYCO) films. They were also found to have moderate dielectric constants (~25 compared to ~277 for $SrTiO_3$). This study was undertaken to further explore the use of $LaGaO_3$, $NdGaO_3$, $SrTiO_3$, MgO, and Y-stabilized ZrO_2 substrates, prepared from single-crystal boules grown by several suppliers using the Czochralski method. Films were prepared by cylindrical magnetron sputtering and laser ablation. Substrate evaluations included measurement of dielectric constant and loss, thermal expansion, and mechanical hardness and toughness. In addition to their moderate dielectric constants, they were found to have satisfactory mechanical properties, except for the twinning tendency of $LaGaO_3$. Lattice mismatch strains were calculated for orthorhombic BYCO films on a number of substrates. $NdGaO_3$ was found to have the best lattice match with BYCO, and is now available twin-free.

Volume 35, Number 1/2, 1991

Preface *by E. J. Farrell, p. 3.*

FEMvis: An interactive visualization tool for mechanical analysis *by G. P. Bala, p. 4.* Increasing competitiveness in the development and manufacturing of mechanically based products requires ever-increasing design and development cycles, and has stimulated the introduction of many computer-based tools to assist with mechanical design and analysis. Currently, strong emphasis is placed upon the *effectiveness* of such tools, and on their enhancement through improved *usability*. A major contribution to improved usability is the level of *interactiveness* of a tool. FEM*vis* is a tool that provides capabilities for interactive visualization of mechanical engineering analysis, including rotation, translation, and magnification of images; views of shape deformations, their time-evolution, and their superposition; visualizations of scalar fields in two and three dimensions using colored isolevels; blending of shape deformation images and isolevel images; visualizations of three-dimensional phenomena by moving a slicing plane through the image, showing cross-sectional deformations and isolevels; and visualizations of multiple shape deformations and multiple scalar fields during a single usage session. FEM*vis* has been implemented in a portable language and a portable graphics package, and can run on a spectrum of hardware platforms from workstations to mainframes. It has been applied to the mechanical analysis of direct-access storage devices (DASD), including stress, strain, modal, and deformation analyses. The interactive nature of FEM*vis* facilitates iterative design refinement and rapid prototyping.

Volume visualization of 3D finite element method results *by K. Koyamada and T. Nishio, p. 12.* This paper describes a method for visualizing the output data set of a 3D finite element method result. A linear tetrahedral element is used as a primitive for the visualization processing, and a 3D finite element model is subdivided into a set of these primitives, which are generated at every solid element. With these primitives, isosurfaces are visualized semitransparently from scalar data at each node point. Two methods are developed for the visualization of isosurfaces with and without intermediate geometries. The methods are applied to output data sets from some simulation results of a semiconductor chip. These are visualized, and the effectiveness of the method is discussed.

Visual interpretation of multidimensional computations and transistor design *by E. J. Farrell, P. A. Appino, D. P. Foty, and T. D. Linton, Jr., p. 26.* As digital simulations and computations become more complex, larger volumes of output are generated; the engineer must select a concise method of displaying the output and extracting relevant information. In this paper we describe an experimental system called the Visual Interpretation System (VIS), which provides a wide range of tools for managing the visualization of simulation data. The effectiveness of VIS results from an interactive interface which controls a database manager and a visualization manager. The database consists of entities called *data sets* that carry a complete description of the geometry and time-dependent behavior of various properties of a simulated physical object. Visualization management involves both 2D and 3D imaging in multiple display windows and animation. Three-dimensional data imaging is based on optical modeling with back-to-front perspective projection. The optical model assigns color and attenuation to each point on the basis of its data value. With the appropriate choice of attenuation and color, the user can display multiple 3D regions, either as solids or transparently. This approach is not based on surfaces, nor does it require the data to have spatial continuity. The usefulness of VIS is demonstrated with data from a large-scale simulation of a transistor. We demonstrate how 3D visualization techniques provide insight into the physics of isolation-trench-bounded devices at both room and low temperatures, which facilitates the development of improved designs.

Data visualization using a general-purpose renderer *by A. Doi, M. Aono, N. Urano, and K. Sugimoto, p. 45.* This paper describes a general approach to data visualization, based on the Rendering Subroutine Package (RSP). (RSP) is a general-purpose polygon-based renderer, and is IBM's first rendering application programming interface (API) for users who wish to develop their own applications. We present an overview of the system, details of the image synthesis tools, and several examples of the application of RSP to architectural CAD, molecular graphics, and computer tomography.

Interactive analysis of the topology of 4D vector fields *by R. R. Dickinson, p. 59.* Interactive visualization methods are now evolving in response to a need to provide more immediate access to particular features of interest to analysts at particular points in the space and time of their data. This paper focuses on feature extraction methods relevant to the analysis of vector fields. In vector fields, "critical points" are those points at which the vector magnitude passes through zero. The word "topology" is used to describe the interconnection patterns between critical points. Topology is central to the understanding of vector fields. It provides very succinct and precise summary information, and can be used to subdivide large fields into well-defined subregions. In this paper, methods for interactively creating maps of vector-field topology are described. The advantages offered by interactive methods in comparison with automatic methods are also discussed.

Picture processing and three-dimensional visualization of data from scanning tunneling and atomic force microscopy *by E. P. Stoll, p. 67.* We present an overview of the current status of picture processing and three-dimensional visualization of data from scanning tunneling microscopy and related techniques. The topics we cover include the physical

basis of the resolution limit and noise sources in scanning microscopes, the design and restoration filters, and methods of visualizing surface contours and other surface properties by use of shadowing, contour lines, and superimposed colors. Postprocessed images of gold, graphite, biological molecules, the active zone of a laser diode, and silicon illustrate the outstanding quality of these methods.

Displaying morphological and lithological maps: A numerically intensive computing and visualization application *by F. Barberi, R. Bernstein, M. T. Pareschi, and R. Santacroce, p. 78.* Algorithms for evaluating digital terrain models (DEMs) and elevation moments such as slope, aspect, relief, and curvature are discussed. Significant new applications based on the elaboration and display of such data are presented. The results show that the processed data can be used for environmental protection and to identify topography-dependent natural-disaster hazards.

IDB: An image database system *by A. Turtur, F. Prampolini, M. Fantini, R. Guarda, and M. A. Imperato, p. 88.* Specialized software and hardware tools are needed to work on digital color images; the usability of an image system implies the availability of such resources inside a coherent environment and a friendly user interface. Furthermore, a large volume of data must be efficiently stored and retrieved. To cope with these problems, the prototype of an image database system, named IDB, has been developed to manage image data in an integrated way. The important features of the system are distributed functions, a multi-user environment, interactivity, and modularity.

Visualizing structure in high-dimensional multivariate data *by F. W. Young and P. Rheingans, p. 97.* We present and discuss several dynamic statistical graphics tools designed to help the data analyst visually discover and formulate hypotheses about the structure of multivariate data. All tools are based on the notion of the "data space," a representation of multivariate data as a high-dimensional (hD) space which has a dimension for each variable (column of the data) and a point for each case (row of the data). The data space is projected orthogonally onto the "visual space," a three-dimensional space which is seen and manipulated by the data analyst. The visual space has a point-like object for each case and can have a vector-like object for each variable. The three dimensions of the visual space are orthogonal linear combinations of the variables. We discuss the notion of a "guided tour" of multivariate data space, and present guided-tour tools, including 1) 6D-rotation, a tool for dynamically rotating, in six-dimensional (6D) space, from one 3D portion of the data space to another while displaying the dynamically changing projection in the visual space; 2) hD-residualization, a tool that determines, at the user's request, the largest invisible 3D space—i.e., the largest 3D space is orthogonal to the visual space. This space is used with the visual space so that 6D-rotation can occur between two new 3D portions of the data space; 3) projection-cueing, a group of three tools that use change in object brightness as a cue to show change in aspects of the projection of objects from the data space to the visual space during hD-rotation. In addition to these tools for touring high-dimensional multivariate space, we discuss tools for manipulating the 3D visual space, and a tool for examining the relationship between two data spaces. Finally, we present a guided-tour implementation in which the user manipulates joysticks and sliders to dynamically and smoothly control the graphics tools in real time. A video supplement demonstrates the implementation.

Visualization of molecular dynamics via ray-tracing and animation in a vectorized environment *by G. N. Williams, E. L. Nelson, D. M. Barnett, and K. Parmley, p. 108.* Scientific visualization methodologies are being utilized increasingly in attempts to understand physical phenomena via mathematical and simulation model results. Presented herein are the results of a visualization project which produced a vectorized, high-resolution, ray-traced animation of the dynamics of a protein molecule. The resulting animation was recorded on 35-mm file, with a resolution of 1024×1024 pixels with 24 color bits. Run-time statistics were also collected which relate image generation parameter ranges and interdependencies.

Animation of computer simulations of two-dimensional turbulence and three-dimensional flows *by M. Briscolini and P. Santangelo, p. 119.* One of the most challenging problems in fluid dynamics is understanding the properties of turbulent flows. The advent of large supercomputers permits the investigation of turbulence with great accuracy in two dimensions, but full three-dimensional problems are physically more complex and their study is currently limited to the case of simple flows. It is shown that the availability of a continuous time-dependent representation of the dynamics of fluid flows can quickly lead to more complete understanding of the many concurrent physical mechanisms ruling turbulence. Some significant examples show how an analog videotape, obtained from direct computer simulations of fluid flows, suggests physical results that can later be obtained through a mathematical analysis of the numerical simulations.

A numerically intensive computing environment: IBM 3090 and the PS/2 Model 80 *by R. F. Arnold, P. Halpern, G. R. Hogsett, B. T. Straka, C. Arasmith, and J. McElroy, p. 140.* Recent advances in personal computer workstations, such as the IBM Personal System/2® Model 80 with its increased memory and CPU speed, loosely coupled with a host IBM 3090™ Processor, can provide considerable computing advantages for executing and visualizing numerically intensive computing (NIC) applications. We have developed a prototype visualization environment which demonstrates effective use of this hardware. The user interface for the NIC application is written using Microsoft Windows® on the PS/2® Model 80 running DOS 3.3. The PS/2 Model 80 is connected to a host 3090 via a PC network. The user enters requests which are application parameters and

selects graphic views for displaying the output results file. The entries are made through user dialog screens on the workstations. The user view of the system is such that it appears that it is running on the workstation. To achieve this transparency, caches are used on both the workstation and the host. The cache on the host is in the form of graphic metafiles and numeric data. The cache on the workstation contains metafiles. Requests are monitored on the workstation to determine whether the results are in the local cache. When they are not, a request file is transferred to the host and checked against the host cache. The NIC application is run only when the requested result is not in either cache. In order to reduce the file size, the results file is converted to a metafile before being transferred to the workstation.

Application of visualization tools in solid mechanics *by S. Moini, p. 156.* With increasingly complex digital simulations and computations, large volumes of numerical output are generated, and users must select more effective techniques for handling and displaying such output in order to extract relevant information. In this study, visualization techniques such as animation, tracking, and 2D/3D color displays are imbedded in implicit and explicit finite element codes for solving complex solid-mechanics problems. With these techniques, the investigator can more fully utilize computer time and better understand the results of long and costly computations. This investigation demonstrates the effectiveness of different visualization techniques and distributed computing on an IBM platform.

Graphic workstations and supercomputers: An integrated environment for simulation of fluid dynamics problems *by F. Piccolo, V. Zecca, A. Grimaudo, and C. Loiodice, p. 167.* An integrated environment for simulation and visualization of physics and engineering problems of industrial interest has been set up at the IBM European Center for Scientific and Engineering Computing (ECSEC). This paper describes the environment, its components, and some experiments carried on at ECSEC to represent 3D objects displayed with the shading technique and the solution of fluid dynamics problems, all treated with the finite element method. Moreover, the paper describes the animation experiments developed to represent dynamics phenomena (fluid flows) and presents a videotape showing the time evolution of three fluid dynamics study cases.

Correlative visualizaton techniques for multidimensional data *by L. A. Treinish and C. Goettsche, p. 184.* Critical to the understanding of data is the ability to provide pictorial or visual representations of those data, particularly in support of correlative data analysis. Despite the many advances in visualization techniques for scientific data over the last several years, there are still significant problems in bringing today's hardware and software technology into the hands of the typical scientist. For example, there are computer science domains other than computer graphics, such as data management, that are required to make visualization effective. Well-defined, flexible mechanisms for data access and management must be combined with rendering algorithms, data transformations, etc. to form a generic visualization pipeline. A generalized approach to data visualization is critical for the correlative analysis of distinct, complex, multidimensional data sets in the space and earth sciences. Different classes of data representation techniques must be used within such a framework, which can range from simple, static two- and three-dimensional line plots to animation, surface rendering, and volumetric imaging. Static examples of actual data analyses will illustrate the importance of an effective pipeline in a data visualization system.

Interactive Quantitative Visualization *by R. L. Peskin, S. S. Walther, A. M. Froncioni, and T. I. Boubez, p. 205.* Interactive Quantitative Visualization, a methodology to enhance scientific and engineering computational simulation prototyping, is defined. Appropriate strategies for implementing IQV in a workstation-based distributed computing environment are discussed. Object-oriented graphical tools and a new data management technique to support IQV and computational steering are described. Two examples of IQV and computational steering are presented: 1) a system to allow interactive solution and visualization of nonlinear boundary-value problems; and 2) a modeling exercise illustrating how IQV and computational steering are used together to prototype simulation of a complex physical system, namely a flag flapping in the wind.

Picturing randomness on a graphics supercomputer *by C. A. Pickover, p. 227.* This paper provides a light introduction to a simple graphics technique which can be used to represent random data on a graphics supercomputer. The representation, called a "noise-sphere," can be used to detect "bad" random-number generators with little training on the part of the observer. The system uses lighting and shading facilities of 3D extensions to the X-Windows or the PHIGS+ standard. To encourage reader involvement, computational recipes and suggestions for future experiments are included.

An interactive graphic tool to plot the structure of large sparse matrices *by G. V. Paolini and P. Santangelo, p. 231.* Many engineering and scientific problems involve the solution of large sparse linear systems. To determine an optimal solving strategy for such systems, it is essential to understand the large- and small-scale properties of the associated sparse matrices. We present a graphic tool to analyze the sparsity pattern and the numeric structure of these matrices. Through examples, drawn from our practical experience, we demonstrate the effectiveness and the interactive features of the tool. These features include zooming, scrolling in different directions, sorting of rows and/or columns, and selective plotting, according to the values of the matrix coefficients.

Visualization in a VLSI design automation system *by D. L. DeMaris, p. 238.* Problems unique to the visualization of

successfully. The implications of these results for hierarchical networks, a class of networks for interconnecting a highly parallel, shared-memory multiprocessor computer system, are discussed.

Volume 35, Number 3, 1991

Preface *by W. Nohilly, S. Hajek, and R. Hawryluk, p. 306.*

The IBM Enterprise System/9000 Type 9121 air-cooled processor *by S. F. Hajek, p. 307.* The IBM Enterprise System/9000™ Type 9121 air-cooled processor achieves, with the same or reduced physical floor space and power levels, a performance level equal to or greater than those of previous IBM processors. This performance level was attained by a combination of design innovations: a new air-cooled thermal conduction module (TCM), integration of bipolar and CMOS technology in this TCM, the design and implementation of a differential current switch bipolar circuit family, integrated programmable memory subsystem design, and extensive use of VLSI technology. The result of these innovations was a 15-ns-cycle air-cooled machine. The salient features and an overview of the machine design are presented.

Differential current switch—High performance at low power *by E. B. Eichelberger and S. E. Bello, p. 313.* The recent IBM System/390® announcement includes six high-performance air-cooled models that use a low-power variation of emitter-coupled logic (ECL). This new logic family, called differential current switch (DCS), uses differential signal pairs to represent logic signals, and combines two-level cascode logic with dotting to implement a complete set of logic circuits. These DCS circuits are described in detail, and the relative value of the DCS and ECL logic families is discussed extensively.

A 128Kb CMOS static random-access memory *by J. L. Chu, H. R. Torabi, and F. J. Towler, p. 321.* This paper describes an all-CMOS 128Kb static random-access memory (SRAM) with emitter-coupled-logic (ECL) I/O compatibility which was designed for the air-cooled Enterprise System/9000™ processors. Access time of 6.5 ns is achieved using 0.5-μm channel length and 1.0-μm minimum geometry. Pipelining and self-resetting circuit techniques permit the chip to operate with cycle time less than access time. To achieve the high-reliability requirement in the TCM environment, a novel technique utilizing a sacrificial substrate is used to "burn in" chips prior to their attachment to the TCM.

IBM System/390 air-cooled alumina thermal conduction module *by J. U. Knickerbocker, G. B. Leung, W. R. Miller, S. P. Young, S. A. Sands, and R. F. Indyk, p. 330.* Advances in multilayer ceramic (MLC) processing, the use of thin-film metallurgy wiring, and enhancements in thermal dissipation, all described in this paper, represent significant milestones in the evolution of microelectronic packaging technology. The IBM System/390® air-cooled alumina thermal conduction module (S/390® alumina TCM) utilizes a 127.5 × 127.5-mm MLC substrate to interconnect as many as 121 VLSI devices and 144 substrate-mounted decoupling capacitors. The substrate provides an array of 648 pads for solder connections to each device, an array of 16 pads for solder connections to each capacitor, and an array of 2772 pins for interconnection with the next package level, and contains approximately 400 m of wiring. The reduced thermal resistance design permits up to 600 W of air-cooling capacity. This paper describes the S/390 alumina TCM fabrication processes and discusses the advances they represent in processing technology, packaging density, and performance. Comparisons to prior technology are made.

IBM Enterprise System/9000 Type 9121 Model 320 air-cooled processor technology *by V. L. Gani, M. C. Graf, R. F. Rizzolo, and W. F. Washburn, p. 342.* The basic component of the new IBM Enterprise System/9000™ Type 9121 Model 320 processor is an air-cooled thermal conduction module (TCM). The fabrication of this module required the integration of new bipolar chips, CMOS SRAM chips, and ECL and DCS logic circuitry in a TCM that could dissipate heat by means of air cooling. The method and details of this process of integration are described and discussed.

An adder design optimized for DCS logic *by A. Weinberger, p. 352.* The basic DCS logic gate provides a two-way SELECT function and, with modifications, a two-way XOR, OR, or AND function. Furthermore, outputs of DCS gates can be wired together (dotted) to perform dotted SELECT, XOR, OR, or AND functions. The versatility of this logic is illustrated in the design of a carry-lookahead adder.

IBM Enterprise System/9000 Type 9121 system controller and memory subsystem design *by B. W. Curran and M. H. Walz, p. 357.* A system controller supporting two processors, two independent memory banks, and a channel subsystem has been implemented within a single air-cooled thermal conduction module for the IBM Enterprise System/9000™ Type 9121 processors. Improvements in technology densities, usage of CMOS and emitter-coupled logic on the same substrate, and innovations in the system controller design were required to achieve the one-module objective. In addition, system reliability is improved with a storage key error-correction code, and storage allocation options are increased with a combined main/expanded store design. In conjunction with the system controller development, a new memory subsystem has been designed for the 9121 system. Innovative large-system memory packaging techniques and functional changes in the data accessing methods have culminated in a memory board which supports up to one-gigabyte system storage.

Design and performance of the IBM Enterprise System/9000 Type 9121 Vector Facility *by T. J. Slegel and R. J. Veracca, p. 367.* The design of the IBM Enterprise System/9000™ Type 9121 Vector Facility is described and its

complex, partially automated design tasks such as VLSI system design are reviewed, and approaches are described. The design domain used to illustrate the approaches is chip-level "floor-planning," an iterative-refinement design methodology for VLSI layout, routing and timing control. The general view structure and control structure are described. Other visualization topics addressed are display of evolving data, sequencing of overlay data, an interleaved temperature–color metaphor for view consistency and clarity, and dynamically generated iconic measurement tools.

Visualizing processes in neural networks *by J. Wejchert and G. Tesauro, p. 244.* A real-time visualization toolkit has been designed to study processes in neural network learning. To date, relatively little attention has been given to visualizing these complex, nonlinear systems. Two new visualization methods are introduced and then applied. One represents synaptic weight data as "bonds" of varying length embedded in the geometrical structure of a network. The other maps the temporal trajectory of the system in a multidimensional configuration space as a two-dimensional diagram. Two-dimensional graphics were found to be sufficient for representing dynamic neural processes. As an application, the visualization tools are linked to simulations of networks learning various Boolean functions. A multiwindow environment allows different aspects of the simulation to be viewed simultaneously using real-time animations. The visualization toolkit can be used in a number of ways: to see how solutions to a particular problem are obtained; to observe how different parameters affect learning dynamics; and to identify the decision stages of learning. A demonstration videotape is provided.

Three-dimensional visualization of many-body system dynamics *by M. Bernaschi, E. Marinari, S. Patarnello, and S. Succi, p. 254.* This paper describes a graphic rendering system for use in visualizing the behavior of three-dimensional physical systems. The tool is general and allows the user to characterize a great variety of phenomena. The only requirement is that the physical system be represented by variables defined on quantifiable positions (sites) within a three-dimensional grid. The variables may be discrete (e.g., binary), real, or even complex numbers. The first part of the paper gives a technical description of the graphic program, which is based on a graPHIGS™ interface; two versions of the code (in the C and FORTRAN languages) are available. The hardware platform consists of an IBM 5080 graphic workstation with a 5081 high-resolution monitor which can be driven either by a machine employing IBM System/370™ architecture with VM/XA™ (in our case a 3090™ processor running under VM/XA) or by a RISC System/6000® workstation [we have used both an IBM RT® System and (recently) an IBM RISC System/6000 processor] running under AIX®. The second part of the paper describes three different examples of the application of this tool: discrete spin models, quantum chromodynamics (QCD), and three-dimensional turbulence. For spin systems and QCD, the physical problem consists in understanding the nature of the phase transition from disordered to ordered states of the system. In both cases a direct (i.e., through visualization) investigation of the system configurations reveals valuable information about properties such as the order of the transition, the behavior of the correlation length, and phase coexistence. We note, however, that the meaning of the site variables is very different in the two cases. In particular, for QCD the site variables are complex numbers, which we code by using a color table to represent the phase of the number and pixel size to represent a value proportional to the modulus. This kind of coding is also used for three-dimensional turbulence. Here the analysis can show where dissipation phenomena take place in the fluid and characterize the geometrical nature of the set of dissipative structures.

Visualizing parallel execution of FORTRAN programs *by F. Szelényi and V. Zecca, p. 270.* As a first step toward the parallel execution analysis of FORTRAN programs, a tool called the Parallel Execution Profiler has been designed and implemented for the graphical postexecution analysis of parallel programs using the Parallel FORTRAN environment as a vehicle for both implementing parallel programs and tracing parallel events. The dynamic behavior of parallel execution is observed interactively in color graphs, which can be displayed concurrently with the source code, and in statistical summaries. This paper describes the implementation of our tool for parallel performance analysis with the aid of a parallelized application program from plasma physics.

A 3072×32-stage TDI imaging device *by E. S. Schlig, p. 283.* A 3072×32-stage TDI charge-coupled imaging device is described. It is believed to be the first reported TDI imager suitable for 300-pel-per-inch document scanning. Its large photocharge capacity gives it the noise performance and dynamic range required for high-quality gray-scale and color imaging in publishing and museum applications. Design options for high-resolution TDI imagers and the uniformity enhancement provided by the TDI mode of operation are discussed.

Traffic studies of unbuffered Delta networks *by P. Heidelberger and P. A. Franaszek, p. 288.* This paper analyzes the performance of unbuffered Delta networks under a nonuniform ("hot-spot") traffic pattern. Particular attention is paid to characterizing the overflow traffic of unsuccessfully transmitted packets. Analytic techniques are used to show that the overflow traffic from an unbuffered packet-switched Delta network is (fractionally) hotter than the offered load. Simulation techniques are used to extend this result to an unbuffered circuit-switched network with limited retrials. In addition, the distribution of the number of trials until a "cold" packet is successfully delivered is shown to have a decreasing hazard rate, which means that it becomes less and less likely with each successive trial that a packet is delivered

performance is evaluated in this paper. The Vector Facility design adheres to the architecture developed for the 3090™ vector facilities. The original design objectives and associated architecture are reviewed. Vector operations and design details are discussed, and specific performance results are shown.

Two approaches to array fault tolerance in the IBM Enterprise System/9000 Type 9121 processor *by P. R. Turgeon, A. R. Steel, and M. R. Charlebois, p. 382.* The system design of the IBM Enterprise System/9000™ Type 9121 processor was intended to provide high performance and dense packaging within an air-cooled system. Packaging and technology factors had a major influence on the fault-tolerance strategies chosen. This paper describes the effect that this design point had on the fault-tolerant capabilities of two critical 9121 array applications. Although the design challenges faced by these array applications initially appeared to be very similar, the resulting solutions represent very different designs with differing fault-tolerance capabilities. The rationale for these approaches is given, and the error-correction algorithms are described.

Enhanced self-test techniques for VLSI systems applied to the IBM Enterprise System/9000 Type 9121 processor *by S. Sarma, p. 390.* This paper discusses the problems associated with obtaining adequate test coverage from random self-test for thermal conduction modules (TCMs) in the air-cooled IBM Enterprise System/9000™ Type 9121 processors. Each 9121 TCM contains approximately a quarter of a million circuits. The present complexity of the TCMs made previous testing methods such as chip-in-place (CIP) testing inviable. The solution was to apply self-test techniques to the 9121 TCMs during the manufacturing process. Analytical and simulation techniques were used to predict the random-pattern testability of the TCMs. The results of the self-test process for the five distinct 9121 processor TCMs are presented. Methods of identifying and modifying random-pattern-resistant logic structures are discussed. It is also proposed that a hybrid approach combining random self-test with deterministic test generation can be used to enhance testability.

Volume 35, Number 4, 1991

Making negatives and plates for printing by electroerosion: Introduction and overview *by K. S. Pennington, p. 458.* The most familiar electroerosion printers operate by removing the whitish or silvery aluminum overlayer from discrete areas of a special paper so as to reveal a black underlayer. The direct negative/direct plate (DNP) material described here is a new dual-function printing material which, when employed together with an electroerosion printer, allows direct generation of negatives or (short-run) offset printing plates suitable for use in printing and publishing applications. The DNP material has been specifically developed for use with the IBM 4250 electroerosion printer family, allowing these printers to generate negatives, plates, and camera-ready copy directly from the computer with no chemical processing.

Making negatives and plates for printing by electroerosion: I. Physical principles *by M. S. Cohen and K. S. Pennington, p. 466.* Electroerosion printing involves removal of the aluminum overlayer from selected areas of a black-coated paper. "Direct negatives" as well as "direct plates" for use in offset lithographic printing may also be generated by electroerosion if a clear polymer sheet is used as the substrate instead of paper, and the black base layer is omitted. If such a substrate is metallized and written by electroerosion, the desired direct negative is created in principle since the metal stops transmitted light and the polyester does not. The direct plate is simultaneously created in principle since the aluminum is hydrophilic and the polyester is hydrophobic. Practical realization of these concepts required studies of the physical principles of the processes involved, which led to techniques for avoidance of mechanical scratching of the aluminum film during writing. For this purpose a mechanically hard underlayer was applied to the substrate under the aluminum, while a very thin lubricating overlayer having some electrical conductivity was applied over the aluminum. The underlayer consisted of silica particles in an organic binder, while the overlayer consisted of graphite particles in a binder. Although scratching is less for smooth than for rough underlayers, rough underlayers were preferred because they offered better writing reproducibility. In particular, debris created during writing was scoured away from the styli in rough-underlayer samples. For writing, a two-phase driver was used, in which the first phase provided a high current for Joule heating with consequent breaking of direct local aluminum–stylus contacts, while the second phase provided an arc which removed the remainder of the aluminum under the stylus.

Making negatives and plates for printing by electroerosion: II. Larger-scale fabrication and testing *by M. S. Cohen, A. Afzali, E. E. Simonyi, M. Desai, and K. S. Pennington, p. 489.* The principles of producing direct negatives and direct plates by electroerosion writing were given in Part I of the present series of papers. Part II is concerned with larger-scale fabrication techniques, characterization methods, and the results of tests, together with their interpretation. Major problems which were encountered are also presented with their solutions. sheet material for the direct negative/direct plate (DNP) was made by 1) Coating polyester rolls with an underlayer. The underlayer coating fluid contained silica, a cellulosic binder, a saturated polyester dispersant, and an isocyanate cross-linker. The coating fluid was carefully milled to control silica particle size before coating. 2) Curing the rolls either at ambient or elevated temperature. 3) Calendering the underlayer. 4) Vacuum-depositing an aluminum film. 5) Coating with an overlayer containing graphite and a binder. Maintenance of good control of material characteristics was

found essential for acceptable functional performance. Among the parameters requiring control were underlayer thickness and surface roughness, aluminum thickness, and overlayer thickness, as well as the composition of the various components. After fabrication, functional testing of the DNP material was carried out on the IBM 4250 printer in order to study the major performance problems of scratching, writing failure, head wear, gouging, and head fouling. Scratching could be suppressed by decreasing the surface roughness and increasing the thickness of the overlayer. Writing efficiency could be improved by increasing the roughness, decreasing the overlayer thickness, decreasing the aluminum thickness, and increasing the pulse length. Head wear could be suppressed by calendering and reducing the roughness. Gouging and excessive head wear could be suppressed by adequate milling, dispersing and filtering of the underlayer coating fluid, and calendering and curing of the coated web, while ensuring good underlayer–substrate adhesion. Fouling was reduced by decreasing the overlayer thickness and reducing the writing pulse length.

Making negatives and plates for printing by electroerosion: III. Use of the direct negative and direct plate *by M. S. Cohen, A. Afzali, E. E. Simonyi, and K. S. Pennington, p. 512.* Issues related to the practical usage in the pressroom of both the direct negative and direct plate are discussed. Two concerns associated with the usage of the *direct negative* during platemaking are treated: 1) the effect of light transmission through defects (voids) in the aluminum film, and 2) the effect of light transmission directly through the aluminum. These concerns may be addressed respectively by careful fabrication of the material and careful control of the exposure conditions during platemaking. Two aspects of the usage of the *direct plate* are considered: 1) the need for a simple prepress "activation" treatment, and 2) press-life limitations caused by wear of the direct plate on the press. Passing a direct plate through an activator solution immediately prior to mounting it on the press prevented scumming in the image area and blinding in the background (nonimage) area. The life on the press of the direct plate is limited by the wearing away of the aluminum from the background areas, thereby causing scumming. Changes in the structure of the direct plate which could prolong press life were investigated.

The art of fractal landscapes *by F. K. Musgrave and B. B. Mandelbrot, p. 535.* Synthetic images of fractal landscapes have moved beyond science to enter the domain of "art for art's sake." We discuss some of the ramifications of this artistic aspect: improving the fractal description of terrains, adding fractal textures to surfaces, and using parallel computers. We illuminate the peculiarities of attaining artistic self-expression in representational imagery purely through formal logic, and discuss its import.

Volume 35, Number 5/6, 1991

Preface *by W. Kleinfelder, p. 571.*

The IBM Victor V256 partitionable multiprocessor *by D. G. Shea, W. W. Wilcke, R. C. Booth, D. H. Brown, Z. D. Christidis, M. E. Giampapa, G. B. Irwin, T. T. Murakami, V. K. Naik, F. T. Tong, P. R. Varker, and D. J. Zukowski, p. 573.* Victor V256 is a partitionable message-passing multiprocessor with 256 processors, designed and in use at the IBM Thomas J. Watson Research Center. Our goals are to explore computer architectures based on the message-passing model and to use these architectures to solve *real* applications. We present the architecture of the Victor system, particularly its partitioning and nonintrusive monitoring. We discuss some of the programming environments on Victor, such as E-kernel, an embedding kernel developed for the support of program mapping and network reconfiguration. We review applications developed and run on Victor and discuss a few in depth, concluding with insights we have gained from this project.

Design choices for the TOP-1 multiprocessor workstation *by S. Shimizu, N. Oba, T. Nakada, M. Ohara, and A. Moriwaki, p. 591.* A snoopy-cache-based multiprocessor workstation called TOP-1 (TOkyo research Parallel processor-1) was developed to evaluate multiprocessor architecture design choices as well as to conduct research on operating systems, compilers, and applications for multiprocessor workstations. TOP-1 is a ten-way multiprocessor using the Intel 80386™ microprocessor chip and the Weitek WTL 1167™ floating-point coprocessor chip. It is currently running under a multiprocessor version of AIX®, which was also developed at the IBM Tokyo Research Laboratory. Our research interest was focused on the design of an effective snoopy cache (all caches monitor all memory-cache traffic) system and the quantitative evaluation of its performance. One of the unique aspects of the TOP-1 design is that the cache supports four different, original snoopy protocols, which may coexist in the system. To evaluate the performance, we implemented a hardware statistics monitor that gathers statistical data. This paper focuses mainly on the TOP-1 cache design—its protocol, and its evaluation by means of the statistics monitor. Besides its cache design, TOP-1 has three other unique architectural features: two independently arbitrated 64-bit buses supported by two snoopy-cache controllers per processor, a communication and interruption mechanism for notifying other processors of asynchronous events, and an efficient arbitration mechanism to allow prioritized quasi-round-robin service with distributed control. These features are also described in detail.

Hierarchically interconnected multiprocessors *by P. A. Franaszek, C. J. Georgiou, and A. N. Tantawi, p. 603.* The

design of interconnection networks is a central problem in parallel computing, especially for shared-memory systems, where network latency, or delay, is one factor that limits system size. This paper discusses aspects of one particular approach to network structure, a design comprising a multiplicity of subnetworks that form a hierarchy of paths. The hierarchy includes fast paths that are used in the absence of contention, and alternate paths with contention resolution. That is, just as in the case of a memory hierarchy, the fastest component of the hierarchy that can provide the desired function is utilized at a given time. The viability and robustness of hierarchical networks is studied first by examining circuit and implementation issues, and then by considering performance modeling and analysis. The overall performance of the hierarchy is shown to be close to that of a contention-free network of fast paths.

Operating system support for parallel programming on RP3 *by R. M. Bryant, H.-Y. Chang, and B. S. Rosenburg, p. 617.* RP3, the Research Parallel Processing Prototype, was a research vehicle for exploring the hardware and software aspects of highly parallel computation. RP3 was a shared-memory machine that was designed to be scalable to 512 processors; a 64-processor machine was in operation from October 1988 through March 1991. A parallel-programming environment based on the Mach operating system was developed, and a variety of programming models were tested on the machine. To help user programs realize the full potential of parallelism on RP3, the RP3 operating system was extended to support such RP3 architectural features as noncoherent caches, local and interleaved storage, and a hardware performance monitor. The system included explicit job-scheduling and processor-allocation facilities, facilities for exploiting the RP3 memory hierarchy, and performance-data collection and logging facilities. This paper describes these components of the RP3 operating system, provides the rationale for the design decisions that were made, and discusses the implementation of these operating system facilities.

The RP3 program visualization environment *by D. N. Kimelman and T. A. Ngo, p. 635.* The performance promised for parallel systems often proves to be somewhat elusive. This paper discusses one important technique for improving the performance of parallel software: *program visualization*—helping programmers visualize the real behavior of an application or system by presenting its state and progress in a continuous graphic fashion. An environment for visualization of program execution is described. Within this *visualization environment*, programmers dynamically establish views of the behavior of a program in execution and watch for trends, anomalies, and correlations as information is displayed. By continually refining the view of the program and replaying the execution of the program, programmers can gain an understanding of program (mis)behavior. This is essential for the debugging, performance analysis, and tuning of parallel software. Design goals for the visualization environment include expandibility, portability, and the ability to accommodate diverse architectures, including highly parallel shared-memory systems and large-scale message-passing systems. Results from visualization of systems and applications running on the RP3, an experimental shared-memory multiprocessor, are presented in the form of color reproductions of typical, useful displays.

The Parallel Processing Compute Server *by E. M. Ammann, R. R. Berbec, G. Bozman, M. Faix, G. A. Goldrian, J. A. Pershing, Jr., J. Ruvolo-Chong, and F. Scholz, p. 653.* The Parallel Processing Compute Server (PPCS) is a distributed-memory multiprocessing system consisting of System/370™ microprocessors (33 at present) interconnected through a matrix switch. This paper describes the hardware configuration, the extensions to the System/370 instruction set that are provided to support the distributed memory and interprocessor signaling, the modifications to the VM/SP operating system that allow it to run effectively on many closely coupled processors (most of which have no disks), and the application-support layer, which permits FORTRAN programs to take advantage of the highly parallel environment. Development of the PPCS is a joint effort of the IBM Böblingen Development Laboratory and the IBM Thomas J. Watson Research Center. A prototype PPCS has been installed at CERN.

Clustering IBM Enterprise System/3090 computers for parallel execution of FORTRAN programs *by L. J. Scarborough, R. G. Scarborough, and S. W. White, p. 667.* Two IBM Enterprise System/3090™ Model 600J computer systems, each with six processors capable of executing vector and scalar instructions, have been connected into a cluster for parallel execution of single FORTRAN programs. The clustering is achieved through a combination of software and hardware. When enabled for parallel execution and allowed to use all twelve processors in the cluster, FORTRAN programs have run as much as 11.7 times faster than when run on a single processor. The combined hardware and software technology is called IBM Clustered FORTRAN. It was achieved by modifying existing technology quickly to provide new capabilities. This paper discusses the modifications and the motivations behind them. It summarizes the performance of several applications executed with Clustered FORTRAN. Finally, it describes how clustering has been used to improve performance in novel ways.

Exploiting database parallelism in a message-passing multiprocessor *by R. A. Lorie, J.-J. Daudenarde, J. W. Stamos, and H. C. Young, p. 681.* Parallel processing may well be the only means of satisfying the long-term performance requirements for database systems: an increase in throughput for transactions and a drastic decrease in response time for complex queries. In this paper, we review various alternatives, and then focus entirely on exploiting parallel-processing configurations in which general-purpose processors communicate only via message passing. In our

configuration, the database is partitioned among the processors. This approach looks promising but offers challenging problems. The paper reports on our solutions to some of them: how to express strategies for efficiently executing complex queries, how to minimize overhead in operations such as parallel joins and sorts, and how to deal with transaction management in a highly distributed system. The paper ends with a discussion of the lessons we learned from exercising a prototype developed in IBM Research.

Multiplication of a symmetric banded matrix by a vector on a vector multiprocessor computer *by R. Reuter, U. Scharffenberger, and J. Schüle, p. 697.* This paper describes how to vectorize and parallelize the multiplication of a symmetric banded matrix by a vector, on a vector multiprocessor. The ideas presented involve two packed-band-storage schemes, and implementations for both schemes are studied. The best among the uniprocessor solutions proposed achieves a maximum of 37.1 Mflops on an IBM Enterprise System/3090™ 400E with Vector Facility (VF). For one of the schemes, a parallel implementation on an IBM 3090™ VF multiprocessor is presented, and time measurements are discussed.

Waveform-relaxation-based circuit simulation on the Victor V256 parallel processor *by T. A. Johnson and D. J. Zukowski, p. 707.* Present-day circuit-analysis tools permit designers to verify performance for circuits consisting of up to 10,000 transistors. However, current designs often exceed several tens of thousands and even hundreds of thousands of transistors. The gap between the number of transistors that can be simulated and the number per design inhibits proper analysis prior to manufacturing, yet incomplete analysis often overlooks design flaws and forces redesign, resulting in increased costs and longer development times. This gap is expected to widen in the foreseeable future. To help close the ever-increasing simulation/design gap, we have developed an experimental parallel circuit simulator, WR_V256, for the Victor V256 distributed-memory parallel processor. WR_V256 has been used to analyze circuits from fewer than 300 to more than 180,000 MOSFETs. WR_V256 was originally based on a Gauss–Seidel relaxation algorithm, which was later replaced with a bounded-chaotic one in order to achieve good parallel speedups for a wider variety of circuits. At this time, speedups of up to 190 have been observed for large circuits.

Further results using the overhead model for parallel systems *by H. P. Flatt, p. 721.* A performance model that takes into consideration the overhead incurred in the use of a parallel system is used to show that the maximum value of the speedup achieved by the parallel system for a fixed problem may be much smaller than the number of processors required to achieve that value. It is also shown that under certain conditions, the problem size may be varied so as to achieve a speedup closely approximating the number of processors used.

The parallel C (pC) programming language *by R. Canetti, L. P. Fertig, S. A. Kravitz, D. Malki, R. Y. Pinter, S. Porat, and A. Teperman, p. 727.* We describe pC (parallel C), an extension of the ANSI C programming language to support medium- to large-grain parallel programming in both shared- and distributed-memory environments. pC aims to make programming for parallel processors accessible to the C community by enriching the C programming model with a small set of constructs supporting parallelism. pC supports shared- and distributed-memory environments via a hierarchical computational model. A pC application comprises a static collection of *tasks* with disjoint memory spaces. A dynamic collection of *threads* runs within each task, sharing the data and code of the task. Language constructs specify concurrent execution of threads within a single task. Additional language constructs specify the interactions between threads through the following mechanisms: initiation of threads in remote tasks by *remote function call*, mailbox-based message passing, and synchronization primitives. The paper introduces the computational model and language constructs of pC and describes a prototype pC compiler and run-time system for the Mach operating system. Several program examples illustrate the utility of pC constructs.

Low-overhead scheduling of nested parallelism *by S. F. Hummel and E. Schonberg, p. 743.* Nested parallelism has the potential not only to permit more parallelism than non-nested parallelism, but to result in better load balancing. However, nested parallelism will not be profitable unless the overhead of scheduling nested parallel constructs can be made nonprohibitive. Previous implementations of nested parallel constructs have been fairly expensive and therefore have not been able to exploit fine-grained nested parallelism. In this paper, we describe a run-time system that schedules a large subset of nested parallel constructs—those that run until completion without blocking—with very little overhead. Our run-time system is built around a novel scheduling policy and work queue. The scheduling policy permits efficient stack-based local-memory storage allocation for task data, which is particularly efficient for multiprocessor architectures with both shared and local memory, such as the RP3. The shared, nonlocking work queue allows processors to obtain tasks in just a few instructions, without sacrificing load balancing.

Execution of automatically parallelized APL programs on RP3 *by W.-M. Ching and D. Ju, p. 767.* We have implemented an experimental APL/C compiler, which accepts ordinary APL programs and produces C programs. We have also implemented a run-time environment that supports the parallel execution of these C programs on the RP3 computer, a shared-memory, 64-way MIMD machine built at the IBM Thomas J. Watson Research Center. The APL/C compiler uses the front end of the APL/370 compiler and imposes the same restrictions, but requires no parallelization directives

from the user. The run-time environment is based on simple synchronization primitives and is implemented using Mach threads. We report the speedups of several compiled programs running on RP3 under the Mach operating system. The current implementation exploits only data parallelism. We discuss the relationship between the style of an APL program and its expected benefit from the automatic parallel execution provided by our compiler.

Automatic partitioning of a program dependence graph into parallel tasks *by V. Sarkar, p. 779.* In this paper, we describe a general interprocedural framework for partitioning a program dependence graph into parallel tasks for execution on a multiprocessor system. Partitioning techniques are necessary to execute a parallel program at the appropriate granularity for a given target multiprocessor. The problem is to determine the best trade-off between parallelism and overhead. It is desirable for the partitioning to be performed automatically, so that the programmer can write a parallel program without being burdened by details of the overhead target multiprocessor, and so that the same parallel program can be made to execute efficiently on different multiprocessors. For each procedure, the partitioning algorithm attempts to minimize the estimated parallel execution time. The estimated parallel execution time reflects a trade-off between parallelism and overhead and is minimized at an optimal intermediate granularity of parallelism. Execution-profiling information is used to obtain accurate execution-time estimates. The partitioning framework has been completely implemented in the PTRAN system at the IBM Thomas J. Watson Research Center. Partitioned parallel programs generated by this prototype system have been executed on the IBM 3090™ and RP3 multiprocessor systems.

Monitoring the performance of commercial T1-rate transmission service *by D. R. Irvin, p. 805.* This paper gathers the scattered empirical and theoretical elements of the performance-management problem for commercial T1-rate transmission service and integrates these elements in a useful way. We propose two variants of a time-based performance-monitoring algorithm that are insensitive to the arrival pattern of transmission errors. The first variant compares a count of errored seconds accumulated over an interval of time to a fixed threshold, and issues an alert to the network operator indicating degraded transmission performance whenever the count exceeds the threshold before the measurement interval expires. The fixed-threshold test is calibrated with reference to the well-known Neyman model of transmission errors on metallic-conductor systems. This calibration is then shown to be suitable as well for monitoring the performance of fiber-optic transmission systems where errored seconds follow the cumulative binomial distribution. The second variant of the new performance-monitoring algorithm replaces the fixed-threshold test with a dual-threshold test having a lower threshold that remains fixed and a higher threshold that floats in response to changes in error characteristics. An analysis based on the difference equations that describe the movement of the floating threshold shows that the dual-threshold test is more responsive than the fixed-threshold test in detecting nonstationary trends toward degraded transmission and in detecting stable but mediocre performance levels.

A trace-driven study of CMS file references *by G. P. Bozman, H. H. Ghannad, and E. D. Weinberger, p. 815.* This paper presents a detailed study of file reference patterns by users of a VM/CMS interactive system. The data were collected from two different IBM locations via CMON, a CMS monitoring facility. We present background information about the CMS file system, the CMON program, and our data-reduction programs, as well as a discussion of the results. Some earlier studies of this type have been restricted to a static analysis of the existing files. However, as is shown in this paper, a static analysis does not reliably reflect dynamic file reference behavior. By using both static statistics and dynamic statistics, it is possible to better understand how file systems are used, to evaluate possible changes, and to provide distribution parameters for modeling. More recent studies of other interactive systems have measured dynamic activity patterns. We compare our results with these when appropriate.